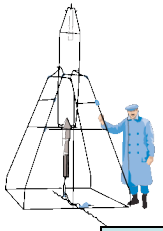

GSFC SEN+HECN Summary Information

J. Patrick (Pat) Gary
Network Projects Leader
Networks and Information Technology Security Group (Code 606.1)
Computational and Information Sciences and Technology Office
NASA Goddard Space Flight Center
Pat.Gary@nasa.gov
September 17, 2009

Presentation for GSFC
Network Evolution and Architecture Transformation
Working Group (NEATWG)

GSFC Managed Networks

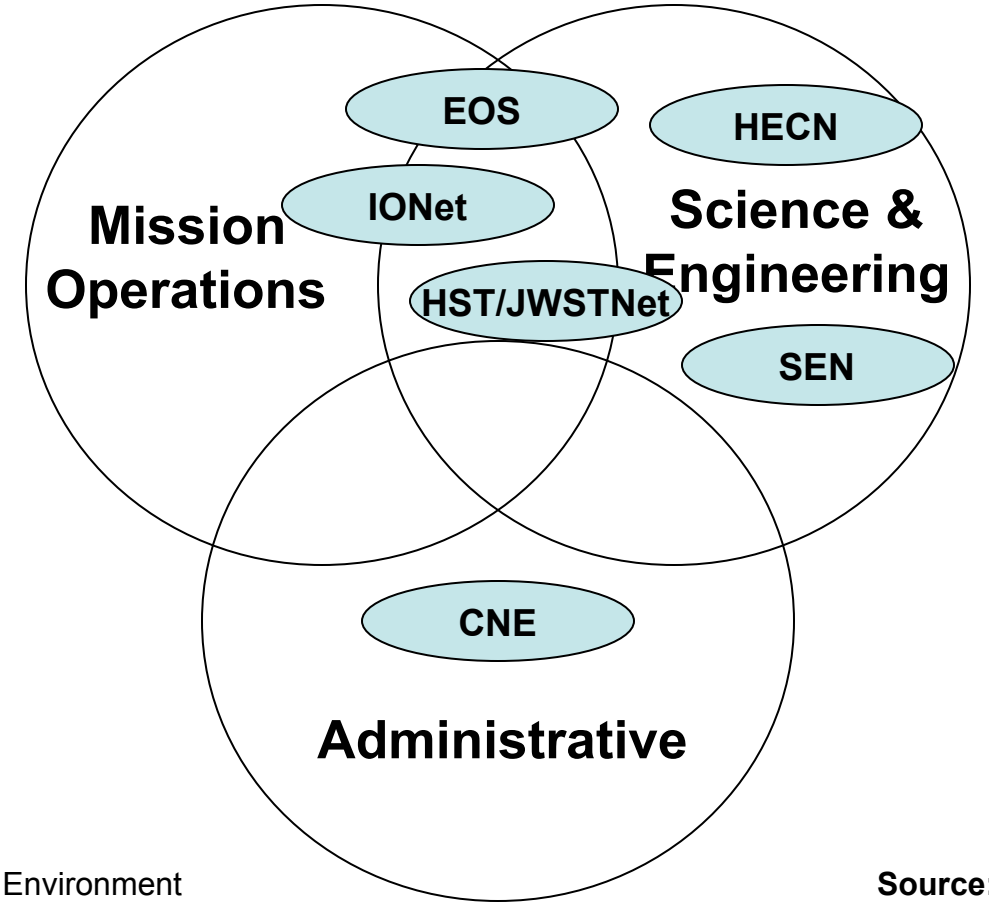

CNE: Center Network Environment
EOS: Earth Observing System
HECN: High End Computing Network
HST/JWSTNet:: Hubble Space Telescope/James Webb Space Telescope Network
IONet: IP Operational Network
SEN: Science & Engineering Network

Source: Craig Hegemann (GSFC)

GSFC SEN+HECN Summary Information

Service Description Overview

- **Scientific and Engineering Network (SEN),** typically enabling 1-10 gigabit per second (Gbps) user connections: a non-mission-dedicated high-end computer network at GSFC serving GSFC projects/users who have computer network performance requirements greater than those baselined for GSFC's general-use campus-wide Center Network Environment (CNE)
  - http://cisto.gsfc.nasa.gov/SEN.html
- **High End Computer Network (HECN)/Lambda Network (L-Net),** typically enabling 10+ Gbps R&D connections: network R&D and testbed evaluations with advanced network technology to contribute to the next generation high-end computer networks at GSFC
  - http://cisto.gsfc.nasa.gov/IRAD_Lambda.html
- Both managed by Code 606.1's HECN Team

Notional Key Characteristics of GSFC's Scientific and Engineering Network (SEN) and High End Computer Network (HECN)/L-Net

[Note: Some data flows/paths are restricted by GSFC security policy and/or management agreement]

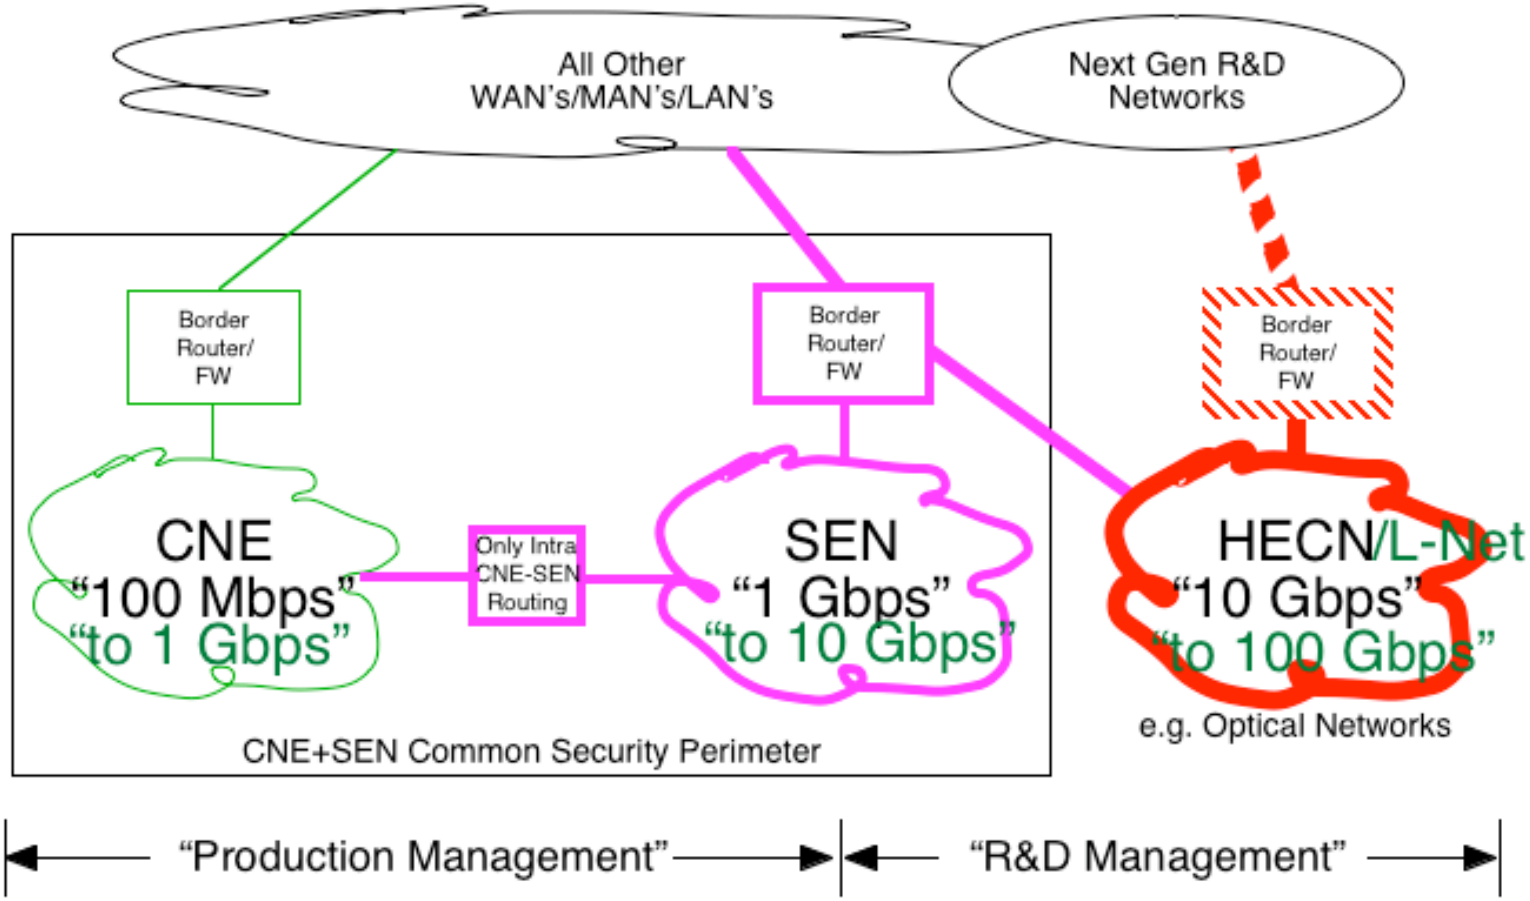

02/27/06

JPG 02/24/04

GSFC SEN+HECN Summary Information

SEN-Connected User Computers (based on 1Apr09 data)

- "Internal Zone"
  - B28: 606.2/NCCS (~110), 606+606.1+606.3 (~20), 610.3/ SVS (~75), 614/SeaWiFS (~12)
  - B32: 586/SAN-testbed (~3), 610.2/DISC (~18), 614.5 (~5), 660.2 (2)
  - B33: 61x (~7)
  - B22+B23: 6xx (~5)
- "Mixed Zone"
  - 606.1/HECN-testbed (~90), 611/ GISS (~240), 614@WFF (0), 614.5/MODIS (~10), 674/CCMC (~10), 694/GGAO (~5)

SEN Connections With Other Networks

- CNE @ 1G & 1G
- EBnet+ IONet @ .6G
- NREN @ 10G
- MAX+Internet2 @ 1G
- DRAGON+NLR @ 3x10G
- NISN L2VPN @ 10G
- NISN PIP @ 10G
- NISN SIP @ 10G

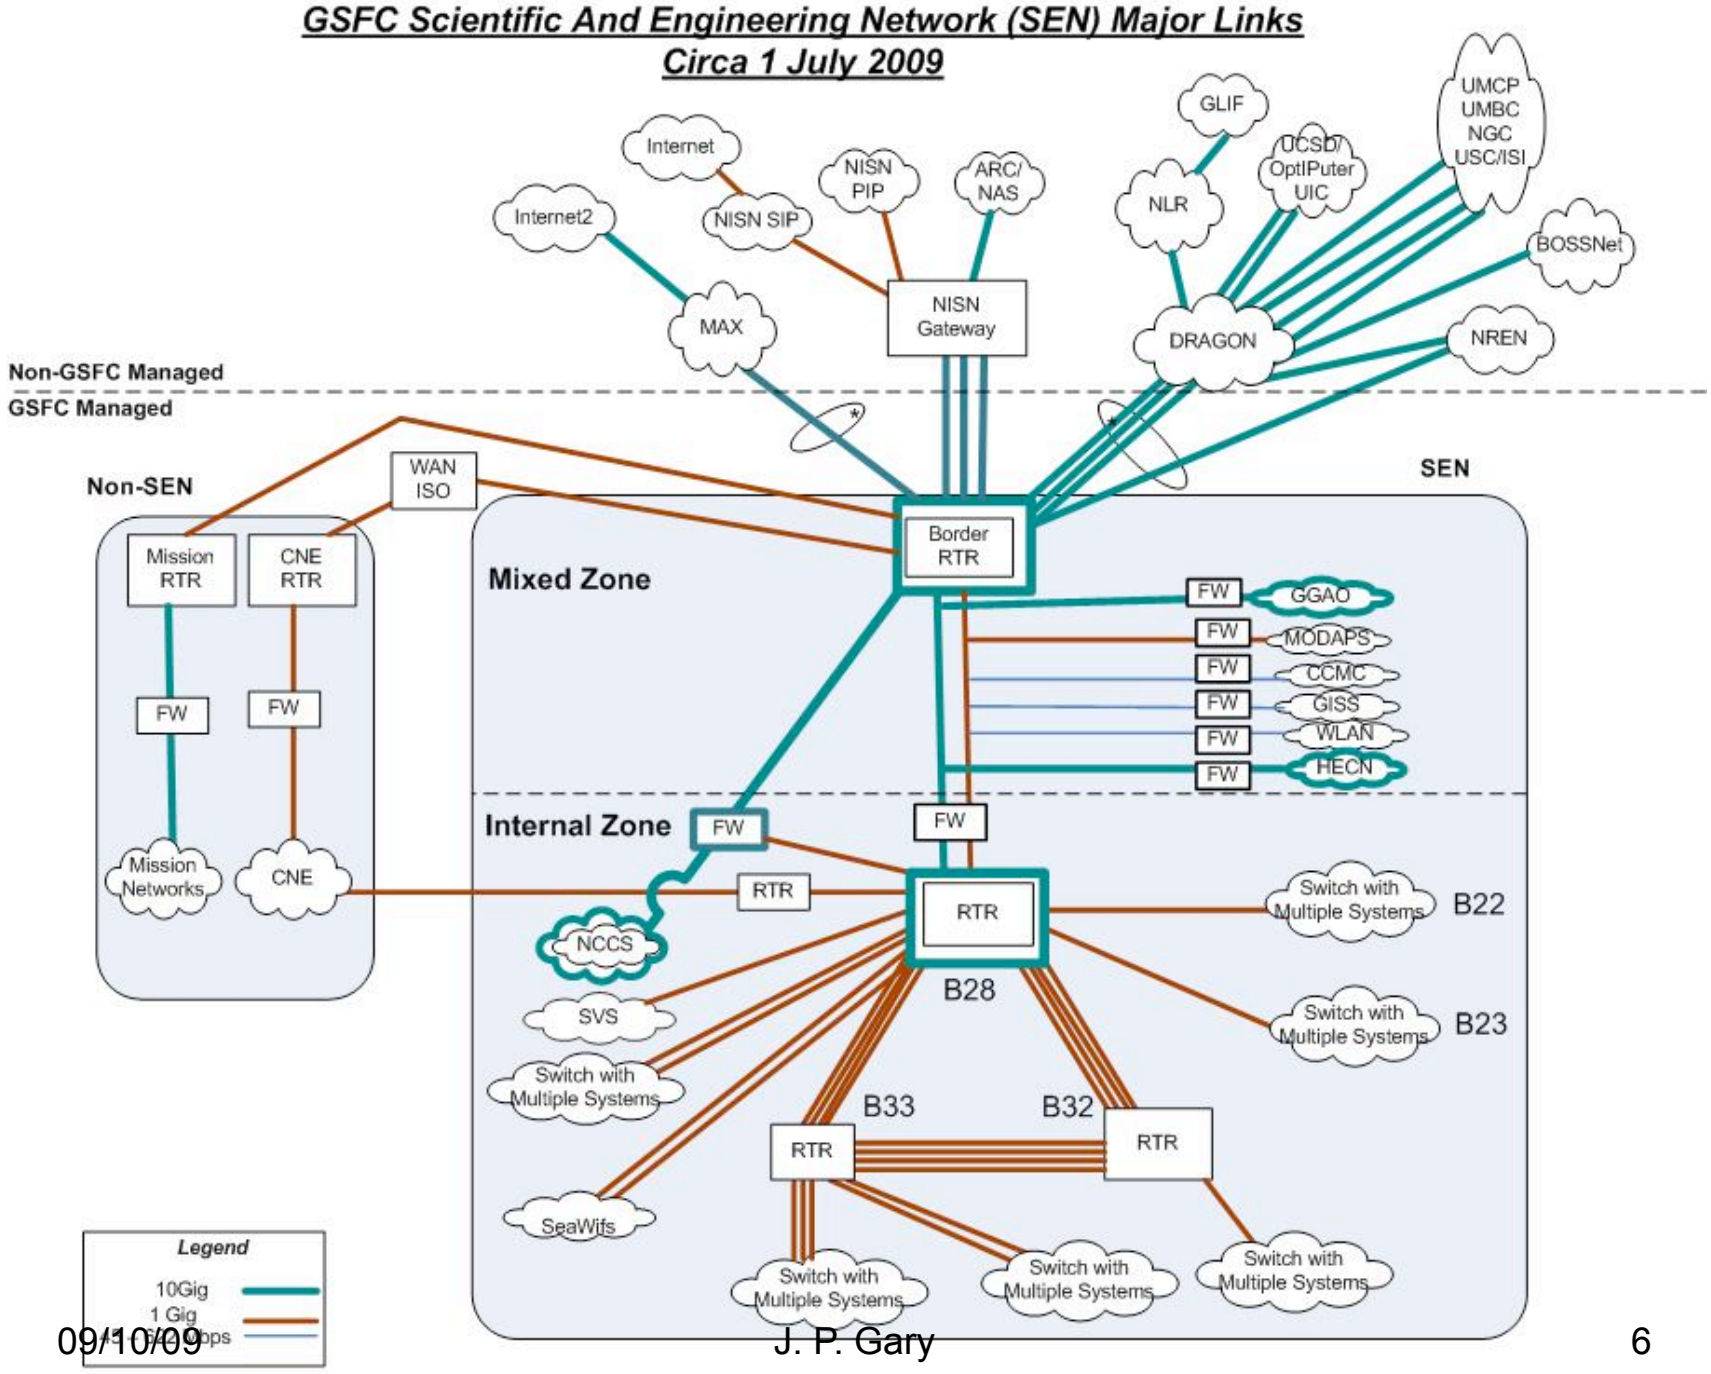
GSFC Scientific And Engineering Network (SEN) Major Links
Circa 1 July 2009
Internet
Internet2
NISN SIP
NISN PIP
ARC/ NAS
MAX
NISN Gateway
GLIF
NLR
UCSD/ OptIPuter UIC
UMCP UMBC NGC USC/ISI
BOSSNet
DRAGON
NREN
Non-GSFC Managed
GSFC Managed
Non-SEN
WAN ISO
SEN
Mission RTR
CNE RTR
FW
Mission Networks
CNE
Border RTR
Mixed Zone
GGAO
MODAPS
CCMC
GISS
WLAN
HECN
Internal Zone
RTR
NCCS
SVS
Switch with Multiple Systems
SeaWifs
B22
B23
B28
B33
B32
Legend
10Gig
1 Gig
* One SEN Physical Fiber Pair and Several Different Lambdas
P.Gary/A.Muppalla
07/15/09

Traffic Analysis for GSFC SEN Interfaces With Other Networks

The statistics were last updated Wednesday, 14 January 2009 at 18:19

`Monthly' Graph (2 Hour Average)

w/CNE: Avg In/Out:23.0/83.6 Mb/s

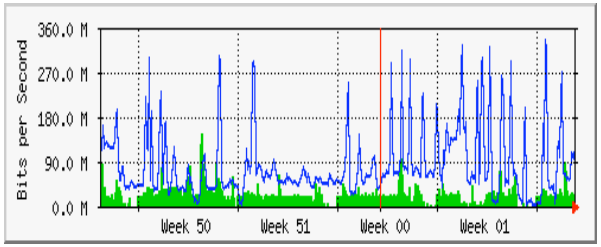

w/NREN: Avg In/Out:634.9/727.6 Mb/s

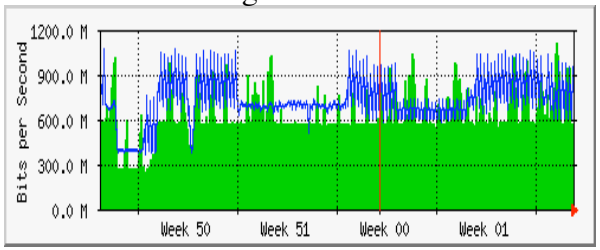

w/L2VPN: Avg In/Out:201.1/403.6 Mb/s

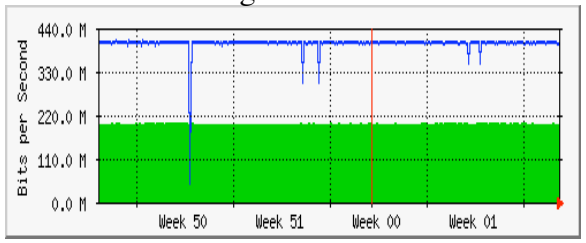

w/EBnet: Avg In/Out:49.4/0.1 Mb/s

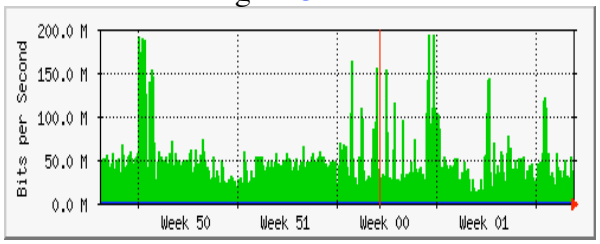

w/MAX+I2: Avg In/Out:27.6/43.2 Mb/s

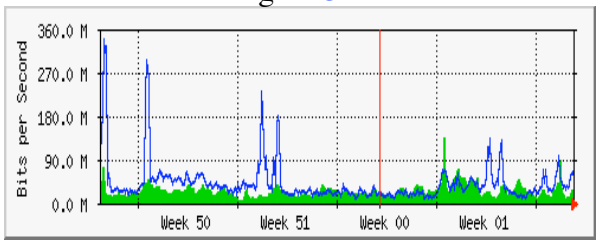

w/PIP: Avg In/Out:0.3/5.6 Mb/s

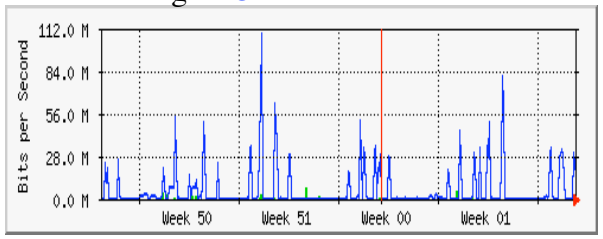

w/DRAGON: Avg In/Out:1.6/1.8 Mb/s

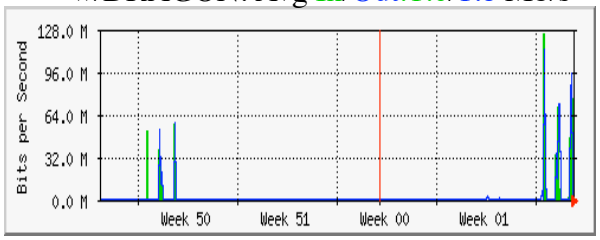

w/SIP: Avg In/Out:0.1/39.9 Mb/s

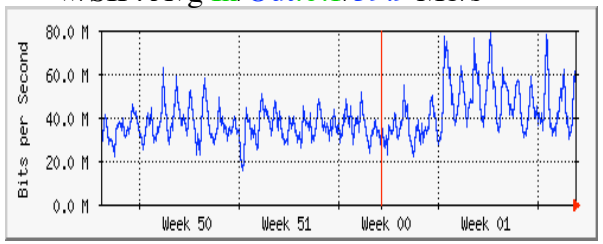

Traffic Analysis for GSFC NREN 6506 10-GigE

The statistics were last updated Wednesday, 26 November 2008 at 11:43

`Daily' Graph (5 Minute Average)

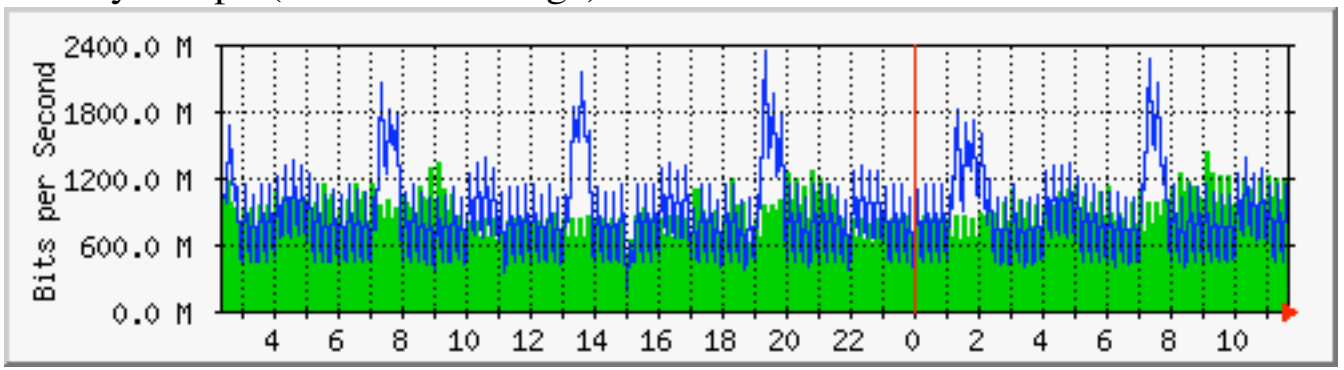

`Monthly' Graph (2 Hour Average)

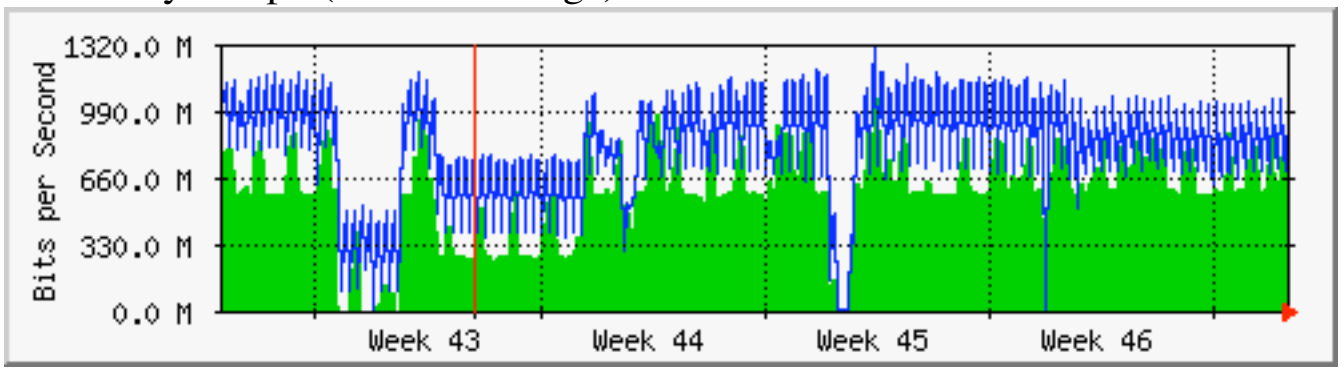

GSFC High End Computer Network (HECN) Project's Research Partners and Collaborators

- *DRAGON Project: http://dragon.maxgigapop.net/twiki/bin/view/DRAGON/WebHome*
  - *PI: Jerry Sobieski (UMCP)*
  - *GSFC L-Net on DRAGON network diagram: http://dragon.maxgigapop.net/twiki/bin/view/DRAGON/Network*
- *e-VLBI Project: http://web.haystack.mit.edu/e-vlbi/evlbi.html*
  - *PI: Alan Whitney (MIT/Haystack)*
  - *GSFC L-Net on e-VLBI network diagram: http://cisto.gsfc.nasa.gov/L-Netpdfs/SC04_eVLBI_network.pdf*
- *GLIF: http://www.glif.is/*
  - *Chair: Kees Neggers (SURFnet)*
  - *GLIF network diagrams: http://www.glif.is/publications/#maps*
- *NGC IT Sector: http://www.it.northropgrumman.com/index.html*
  - *PI: Brice Womack (NGC)*
  - *GSFC L-Net on NGC IT Sector Colshire network diagram: http://cisto.gsfc.nasa.gov/L-Netpdfs/DRAGON_NGC_030606.pdf*
- *NLR: http://www.nlr.net/*
  - *CEO: Tom West (NLR)*
  - *NLR network diagram: http://www.nlr.net/infrastructure/*
- *NREN Project: http://www.nren.nasa.gov/*
  - *PM: Ken Freeman (ARC)*
  - *GSFC L-Net/SEN on NREN network diagram: http://cisto.gsfc.nasa.gov/L-Netpdfs/CENIC2006_13_mfoster_excerpts.pdf*
- *OptIPuter Project: http://www.optiputer.net/*
  - *PI: Larry Smarr (UCSD)*
  - *GSFC L-Net on OptIPuter network diagram: http://cisto.gsfc.nasa.gov/L-Netpdfs/SMARR-OptIPuter-AHM-gold.pdf*
- *TeraFlow Testbed Project: http://www.teraflowtestbed.net/*
  - *PI: Robert Grossman (UIC)*
  - *GSFC L-Net on TeraFlow Testbed network diagram: http://www.ncdm.uic.edu/maps/index.jpeg*

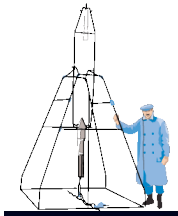

Global Lambda Integrated Facility **World Map**
http://www.glif.is/publications/maps/GLIF_5-08_World_4k.jpg

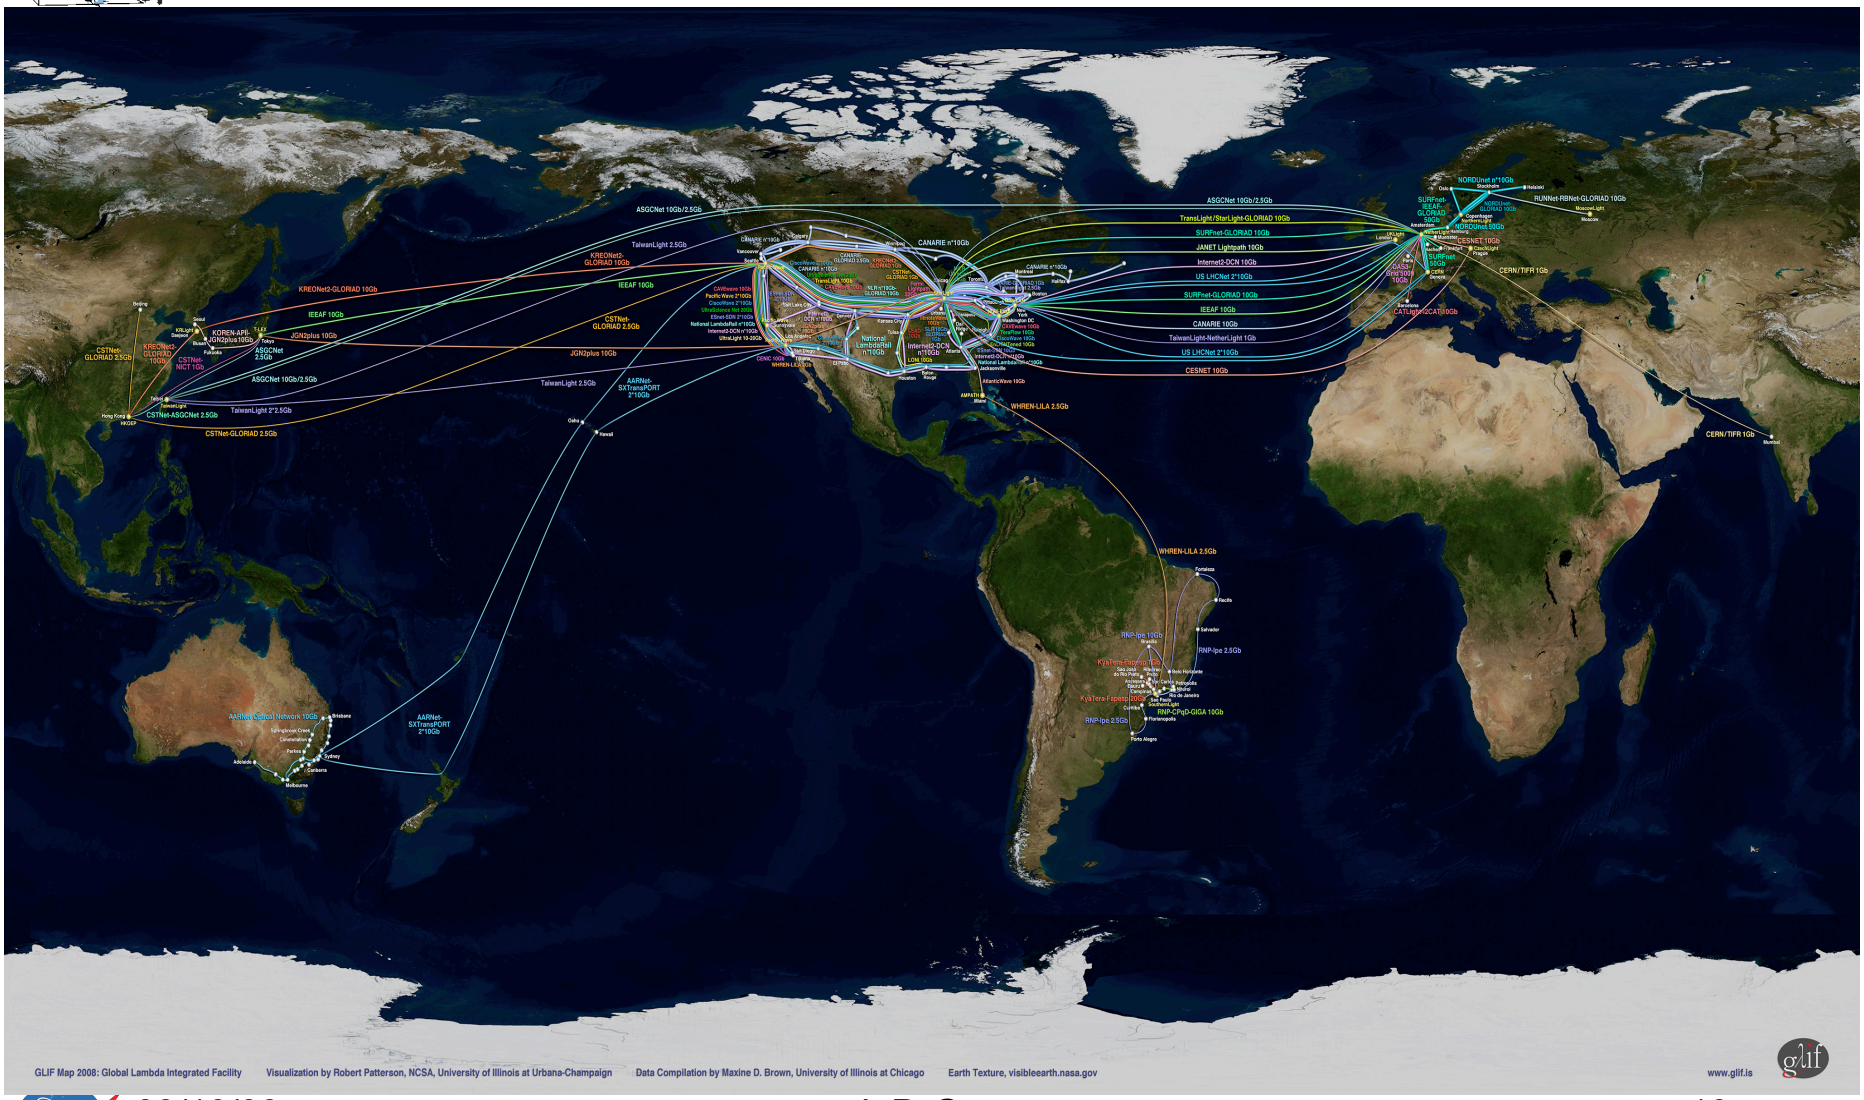

Global Lambda Integrated Facility North America Map
http://www.glif.is/publications/maps/GLIF_5-08_US_4k.jpg

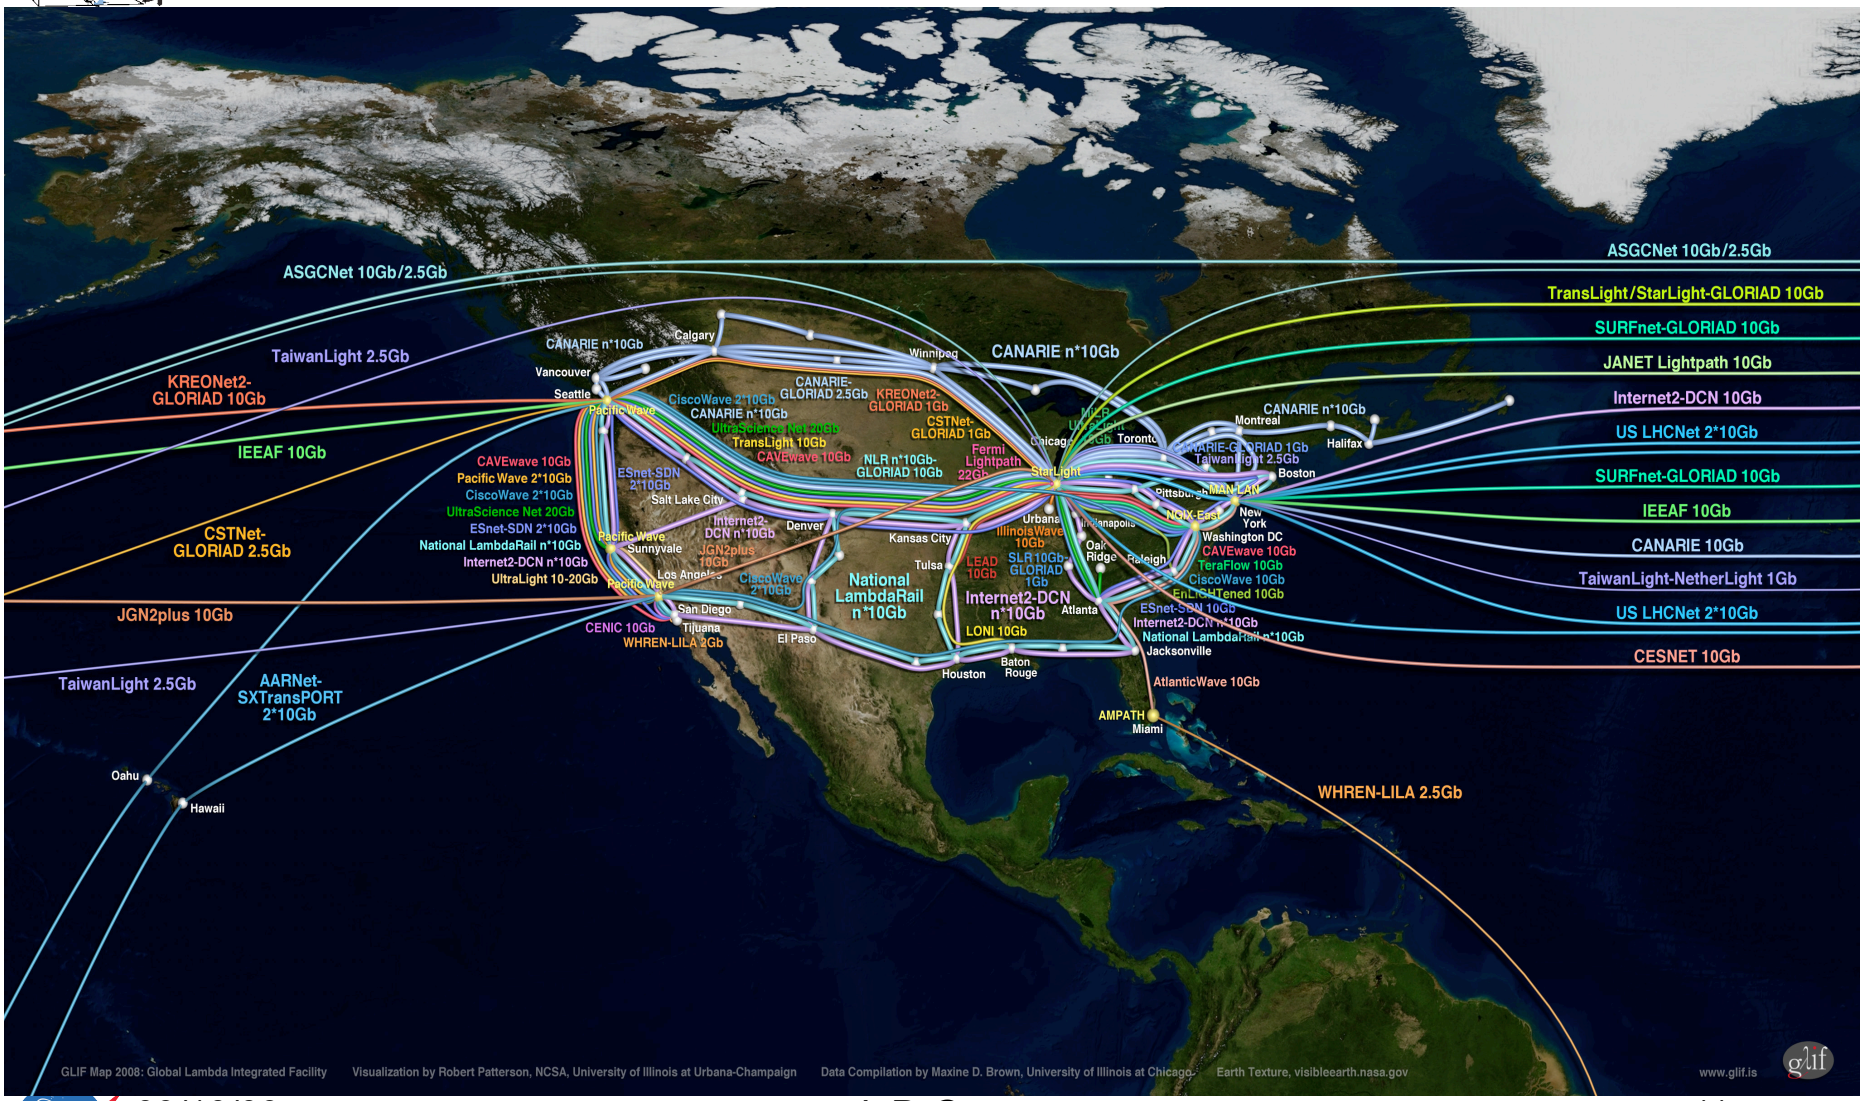

Combining Telepresence with Remote Interactive Analysis of Data Over NLR

www.calit2.net/articles/article.php?id=660

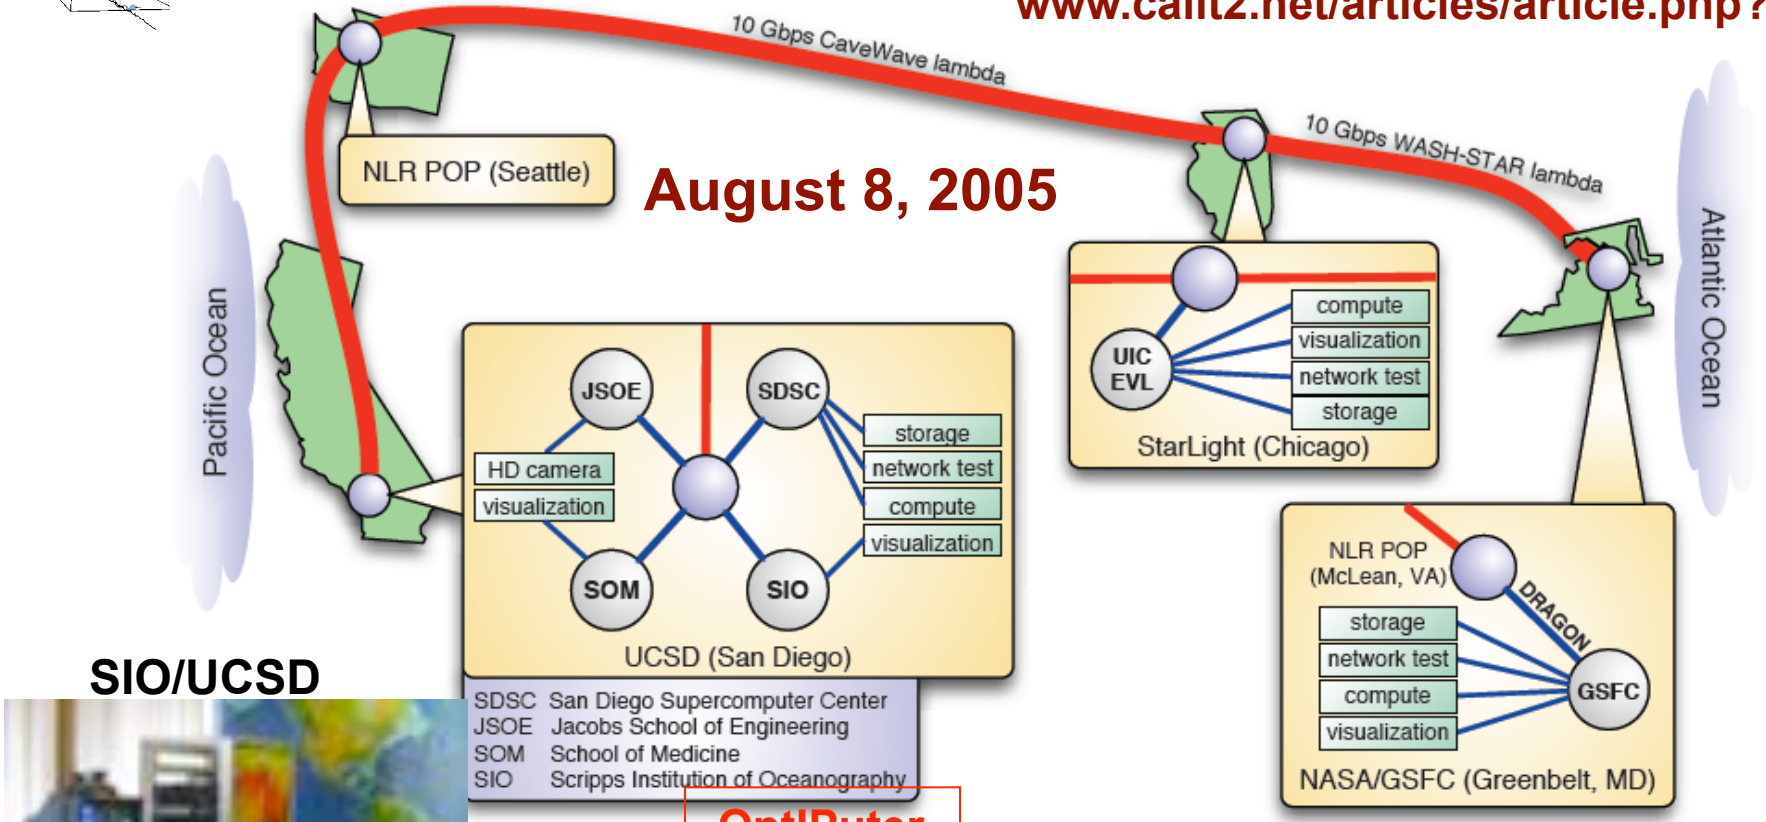

SIO/UCSD

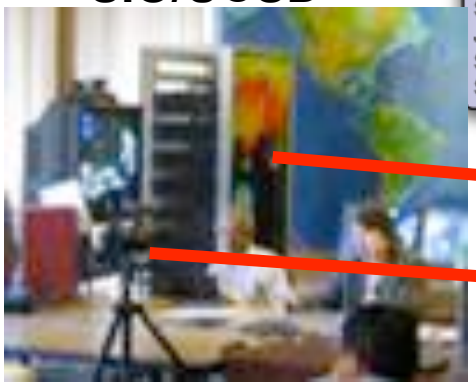

OptIPuter Visualized Data

HDTV Over Lambda

NASA Goddard

GSFC Network Support for DRAGON Xnet Demo During SC06

Streaming NASA HD Video Uncompressed in Realtime from GSFC to the SC2006 Showroom Floor in Tampa

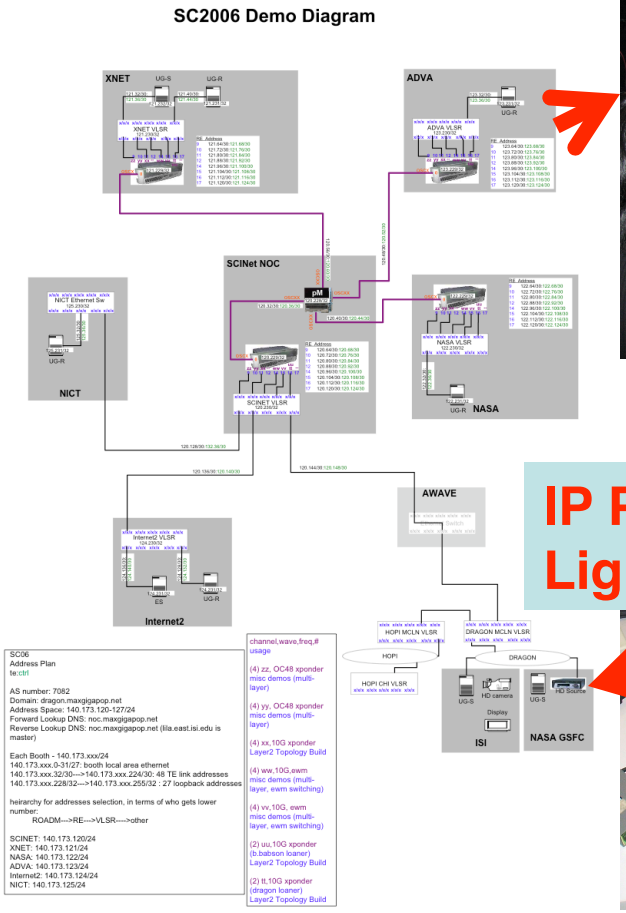

High level network diagram, prepared by DRAGON's Chris Tracy, showing the optical WAN pathways between GSFC and the five booths at SC06 hosting DRAGON's Xnet demo.

At the Internet2 booth at SC06 (one of five hosting DRAGON's Xnet demo) realtime uncompressed HD video from the High End Computer Network (HECN) Team's lab at GSFC is displayed to SC06 attendees.

A Panasonic AJ-HD1200AP HD player, loaned from GSFC's TV Studio (courtesy of Pat Kennedy), provides one of the HD video stream sources.

IP Packets in LightPaths

HD Video

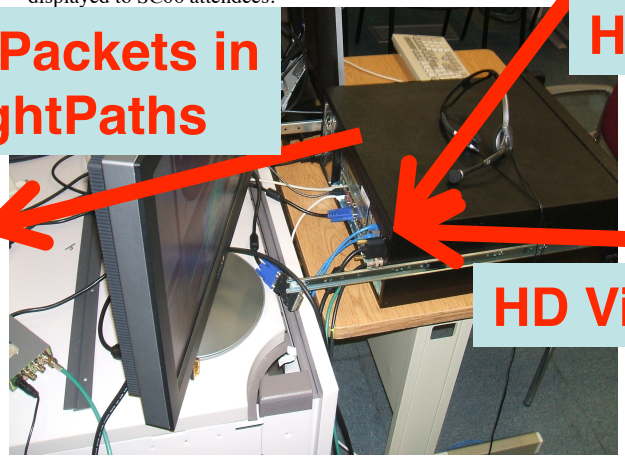

UltraGrid software and a HD video capture/compression card, loaned from USC/ISI-East's Tom Lehman, in HECN's Pentium4 IP-packetizes and transmits the digital video at 1-Gbps through an optical WAN path dynamically provisioned by DRAGON's network control-plane software.

HD Video

A Hitachi SK-3010P HD camera, loaned from GSFC's TV Studio (courtesy of Pat Kennedy), provides one of the realtime HD video stream sources.

GSFC SC07 Highlight: Special Network Support for Real-Time 3D Demo

- ARC and GSFC high-end networking teams and partners established special links for the 3D Holographic Display System in the NASA exhibit. Using a coordinated set of VLANs to carry two separate 1-Gbps flows, the links extended from SC07's local SCinet across the PacWave and NREN to GSFC then back to the exhibit, requiring a 185-millisecond round-trip time.
- The links carried two real-time, uncompressed, high-definition video streams with 1980x1080, 16 bit-pixels per progressive frame from a pair of stereoscopically aligned HDcams mounted above the display.
- Significantly, the links maintained perfect time-synchronization so that the video streams could be combined and displayed as non-distorted holographic images even after their round-trip transmission.

The 3D Holographic Display System is a true 3D visualizer that does not require goggles or other forms of head-gear.

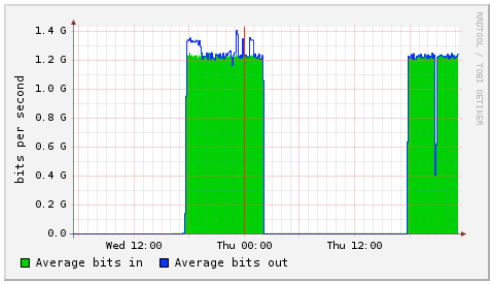

These measurements show the ~1.2 Gbps of bandwidth needed to transport the 22 frames-per-second video streams.

POC: Pat Gary, James.P.Gary@nasa.gov,
(301) 286-9539, CISTO Networks Manager

Source: Dan Magorian (MAX)

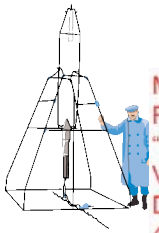

MAX NASA JUNIPER
FUJITSU 40G INTEROP
"Snake Testing" setup
V1.1 8/6/08
Dan Magorian
40g@maxgigapop.net

Color Legend:
Grey: MAX existing
Yellow: Fujitsu loaner
Green: Juniper loaner
Blue: NASA loaner

The approach used behind NASA and MAX's "snake testing" is that because 40G applications are rare and OC768 testers expensive, 10G streams from line-rate bit blasters are muxed together and hard looped between the 10G interfaces to fill the OC768 lambda. MPLS and possibly virtual routers will be used on the Junipers to separate the streams. DWDM pluggable XFP optics will be used on the NASA bit blasters and the College Park Juniper 10G interfaces to enable the blasters to remain at NASA/GSFC for convenience. Juniper loaned the T1600 routers so that there would be no interaction with the MAX production T640 routers with the upgrade to support the 4x 10G pic. Any issues with the Fujitsu 40G test lambdas should not impact the MAX production 10G lambdas over the existing diverse paths.

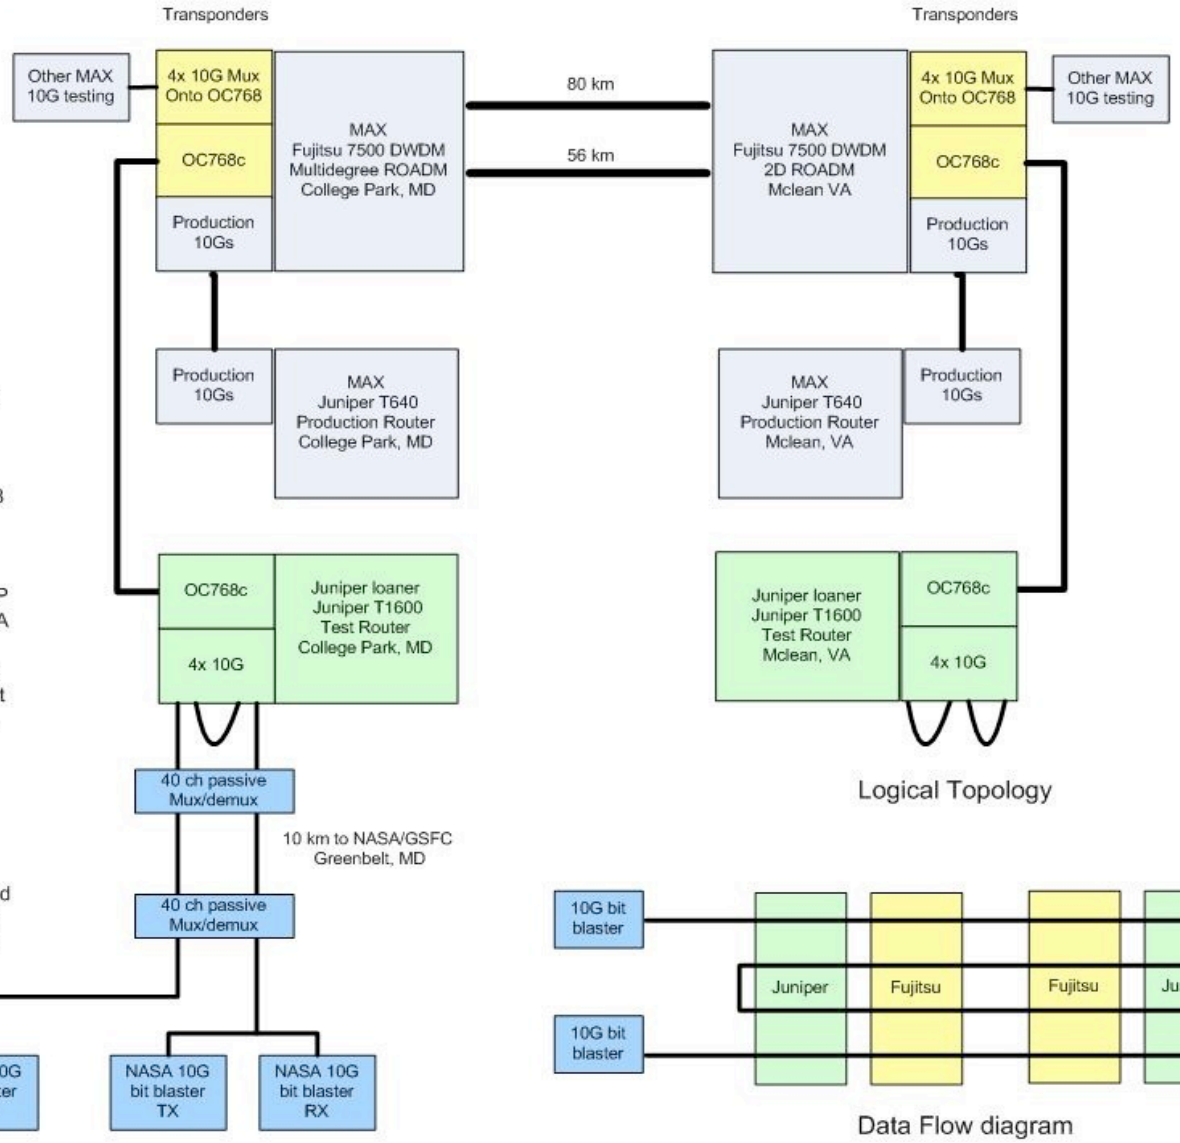

A Sample Traffic Analysis During MAX's 40-Gbps Interface Testing Between CLPK and MCLN

A joint testing effort among Fujitsu, Juniper, MAX & NASA/GSFC

The statistics were last updated **Tuesday, 14 October 2008 at 9:27**

`Daily' Graph (5 Minute Average)

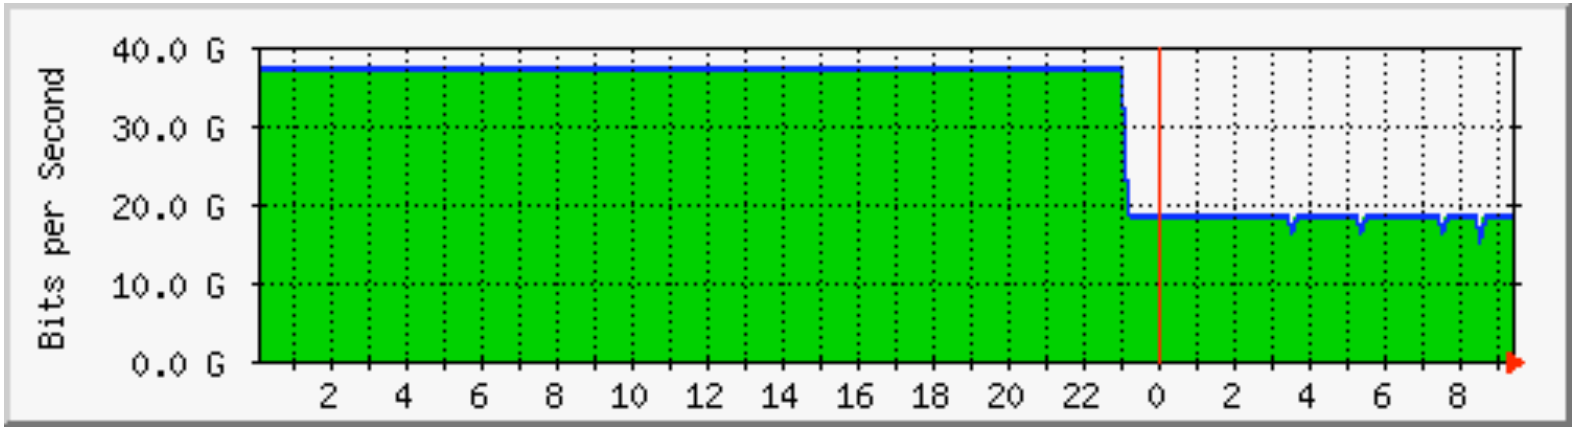

Max In:37.5 Gb/s (94.3%)Average In:31.5 Gb/s (7.8%)Current In:18.7 Gb/s (47.0%)
Max Out:37.6 Gb/s (94.4%)Average Out:31.5 Gb/s (7.8%Current Out:18.7 Gb/s (47.0%)

Snapshot taken during MAX NASA JUNIPER FUJITSU 40G INTEROP Test: One 10G stream from each of two pair of NASA line-rate nuttcp-servers is hard looped between some Juniper T1600 10G interfaces to fill the Juniper OC768c interfaces and Fujitsu 40G optical transponders set up between MAX POP's at College Park, MD and McLean, VA.

Source: Venkatram Vishwanath (UIC)

Striding through remote MERRA dataset at Chicago from GSFC

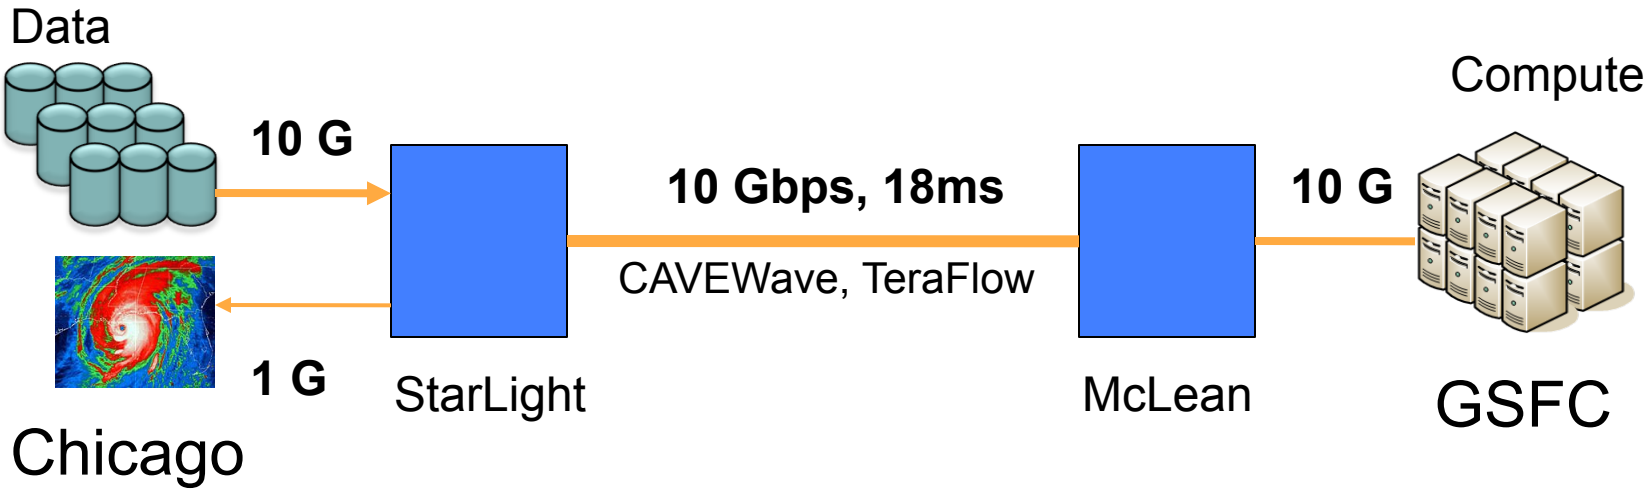

- 10 Gbps over the Teraflow/CAVEWave Optical Network between Chicago and Goddard
- Data located at EVL, Chicago
- Computation at Teraflow Nodes in Goddard
- Visualization at EVL, Chicago

Source: Venkatram Vishwanath (UIC)

Network usage during LambdaRAM experiments 27 October 2008

TenGigabitEthernet 1/1 CAVEwave to DC (NLR-STAR-WASH-10GE-103) -- evl-e600

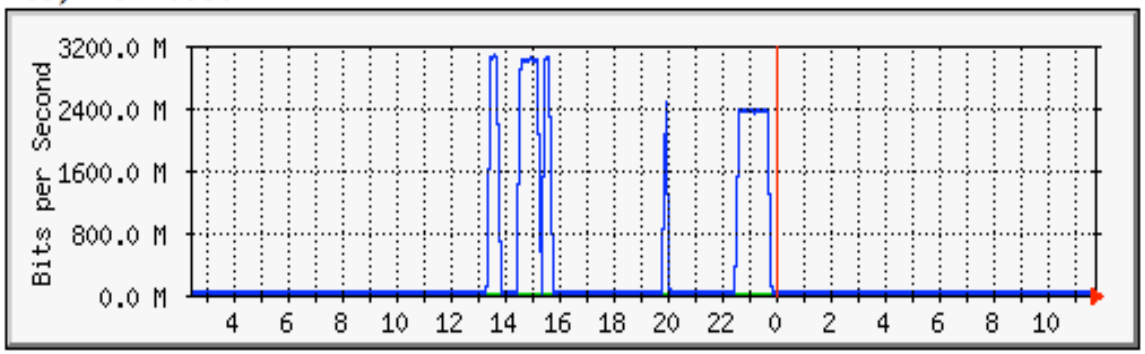

30 min average

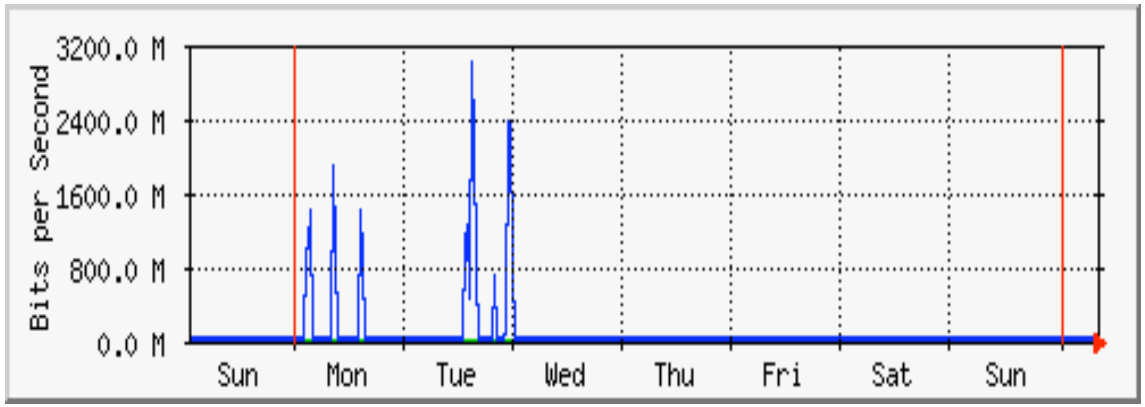

Daily average

Introduction To GSFC High End Computing 20, 40 &100 Gbps Network Testbeds

Problem Statement

- GSFC's NASA Center for Computational Sciences (NCCS) is increasing its data production/analysis/storage capacities and capabilities
- Higher bandwidth networks can be deployed
- But something in the combination of our file copying applications, disk I/O subsystems, server/workstation configurations, protocol stack tuning and/or NICs is preventing full use of our higher bandwidth networks

Introduction To GSFC High End Computing 20, 40 &100 Gbps Network Testbeds

Part of the Solution

- Prepare applicable testbeds to identify the bottlenecks and investigate solutions/alternatives
  - Plan multi-phased LAN, MAN & WAN testbeds with 20, 40 & 100 Gbps throughput performance objectives
  - Assemble low-cost high-performance network-test workstations to assess/troubleshoot:
    - New network technologies
    - File copying application tuning
  - Arrange joint collaboration on WAN file-accessing applications with NCCS and NASA Advanced Supercomputing (NAS) Division data management experts to ensure that the network testbed efforts will have at least HEC Program applicability
- Transition findings into production environments

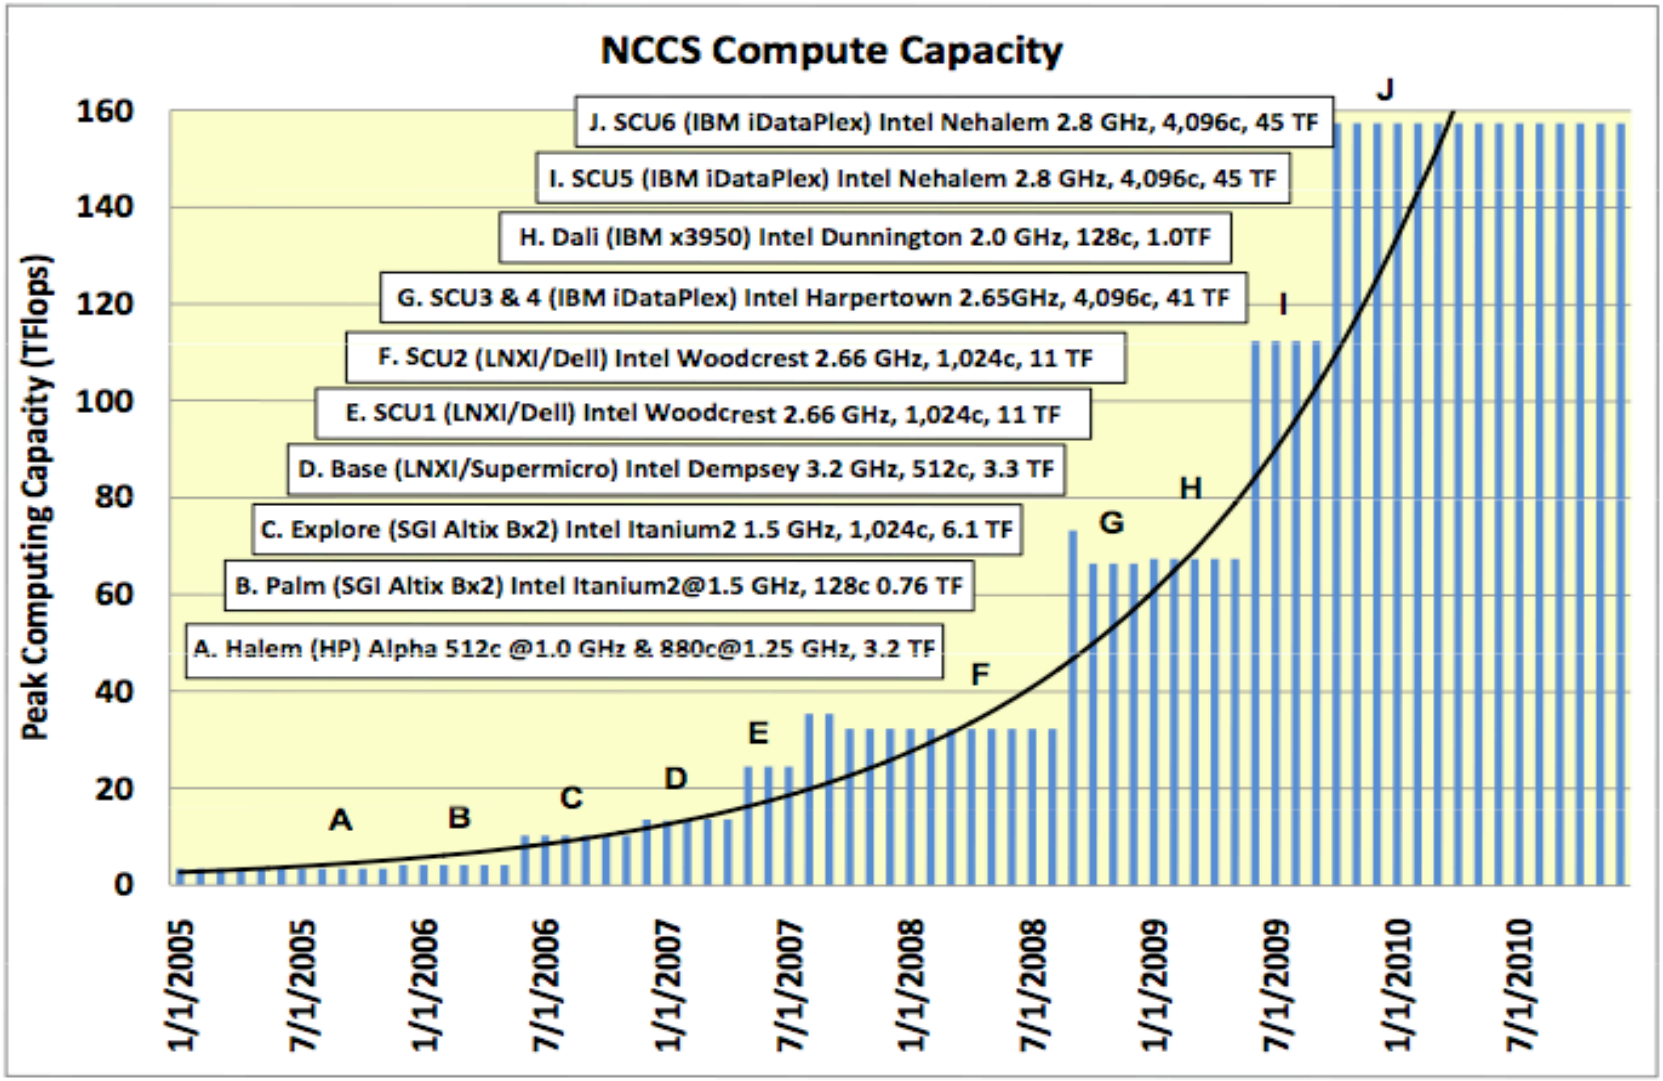

Source: Dan Duffy/GSFC (GSFC/NCCS) & Scott Wallace (CSC) (GSFC/NCCS)

Excerpt from HEC Program Monthly Status Report Apr09
Source: Sally Stemwedel/GS&T (GSFC/NCCS)

NCCS Usage for April 2009

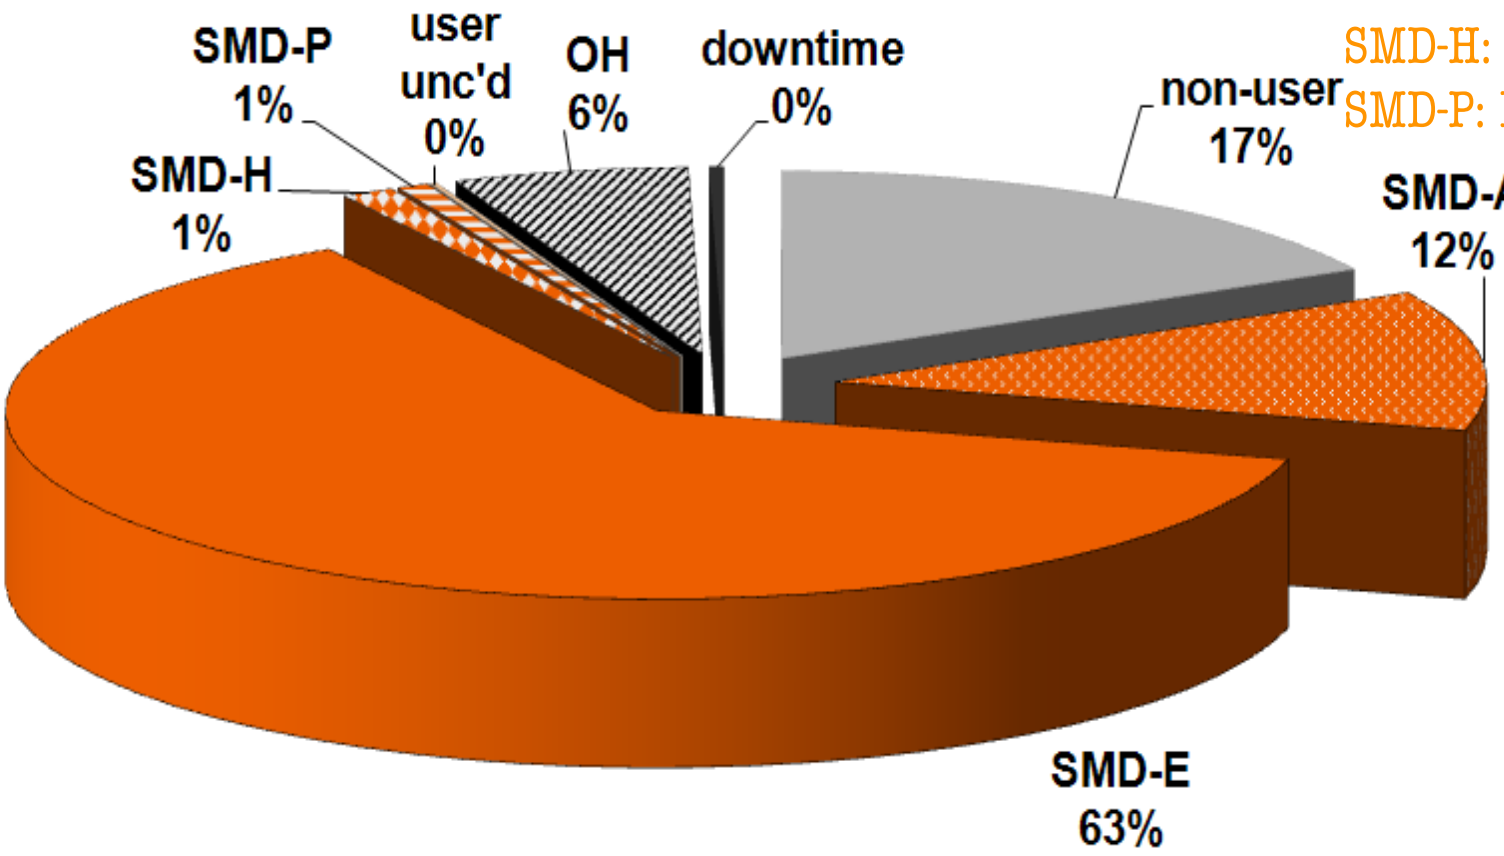

Science Users
SMD-A: Astrophysics
SMD-E: Earth Science
SMD-H: Heliophysics
SMD-P: Planetary

Excerpt from HEC Program Monthly Status Report Apr09
Source: Sally Stemwedel/GS&T (GSFC/NCCS)

SMD PI Organization Distribution

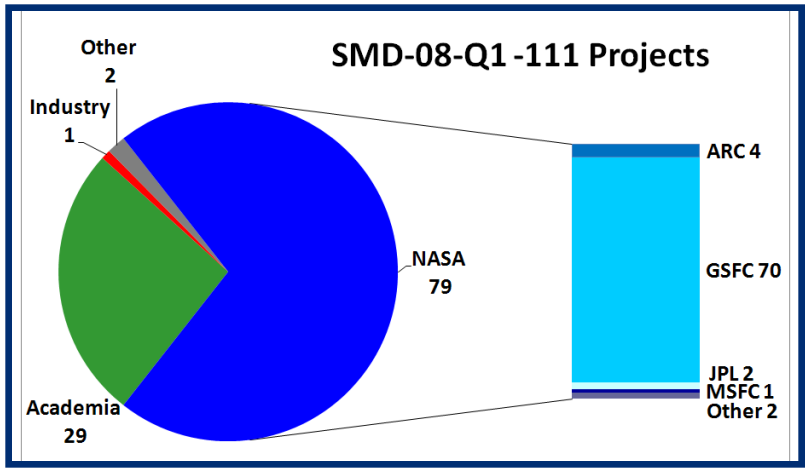

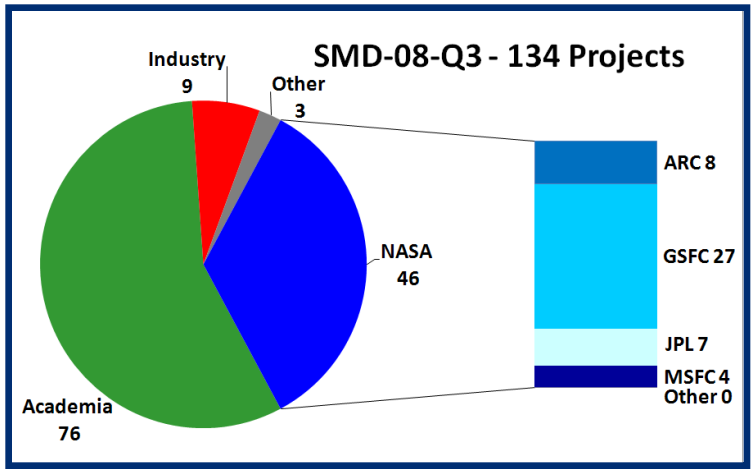

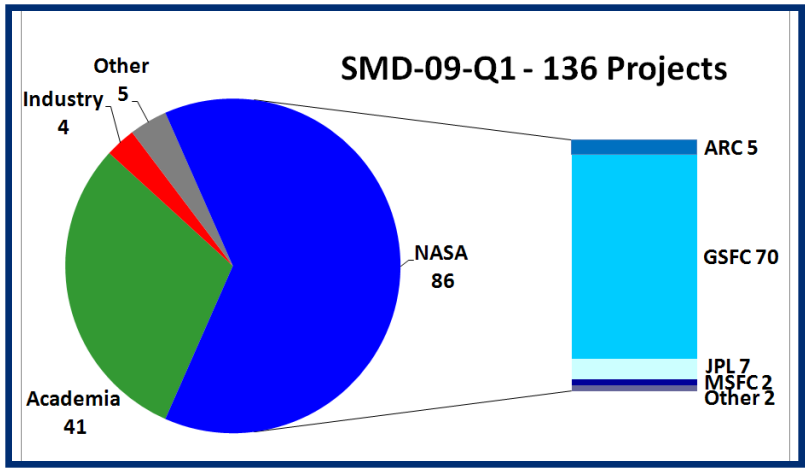

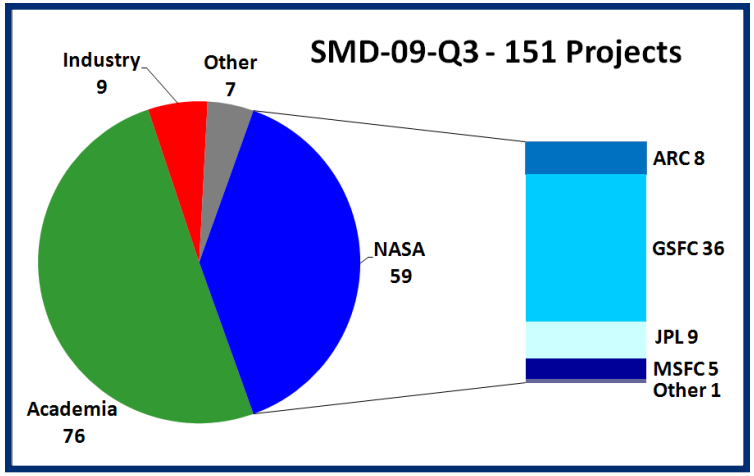

Introduction To GSFC High End Computing 20, 40 &100 Gbps Network Testbeds

Multi-Phased Network Testbed Objectives

- A series of partly overlapping approximately-two-year Phases
- Each producing a demonstrable network capability with end-to-end throughput performance goals targeted respectively at 20, 40 & 100 Gbps
- Each with joint collaboration on WAN file-accessing applications with NCCS and NAS data management experts to ensure that the network testbed efforts will have at least HEC Program applicability

[1] P. Gary, "Plan/Proposal For Initiation of 20, 40 & 100 Gbps Network Technology Testbeds", 29Jan09 draft proposal for potential submission to the NASA Strategic Investment Business Case (SIBC) Initiative for Networking R&D that is expected to emerge soon within NASA.

[2] D. Duffy et al, "NASA High End Computing Selected Wide Area Network Testing: Representative Test Plan", 18Mar09 draft.

Both [1] and [2] were emailed to HEC's network engineers at NAS and NCCS.

Introduction To GSFC High End Computing 20, 40 &100 Gbps Network Testbeds

Example Throughput Performance Benchmarking

- Several representative test types are described in [2], including ones involving:
  - nuttcp at wire transfer speeds for the much needed performance baseline against which all subsequent performance measurements will be compared
  - scp, bbftp, nfs, iRODS, GridFTP and other file transfer mechanisms that NASA HEC users jobs use at both NAS and the NCCS to move data back and forth over the wide area
  - iSER and iSCSI over RDMA to mount the disks either locally or remotely
  - GPFS and/or Lustre shared file system
- Actual throughput performance results will be widely distributed once they are available

Introduction To GSFC High End Computing 20, 40 &100 Gbps Network Testbeds

Phase 1 Network-Test Workstation Functional Objectives with Performance Targets (1 of 2)

- "B" (Baseline) systems:
  - Primarily for network throughput evaluations via nuttcp memory-to-memory testing at up to 40-Gbps unidirectional, 40-Gbps bidirectional (80-Gbps "total")
  - Secondarily for WAN file copying application throughput evaluations in disk-to-disk testing at up to 10-Gbps unidirectional
- "C" systems:
  - Primarily for WAN file copying application throughput evaluations in disk-to-disk testing at up to 20-Gbps unidirectional
- "A" systems:
  - Primarily for WAN delay simulation at up to 40-Gbps unidirectional, 40-Gbps bidirectional (80-Gbps "total")
  - Also as firewall at up to 20-Gbps unidirectional, 20-Gbps bidirectional (40-Gbps "total")

Introduction To GSFC High End Computing 20, 40 &100 Gbps Network Testbeds

Phase 1 Network-Test Workstation Functional Objectives with Performance Targets (2 of 2)

- "A+" systems:
  - Primarily for network throughput evaluations via nuttcp memory-to-memory testing at up to 70-Gbps unidirectional, 40-Gbps bidirectional (80-Gbps "total")
    - Actual performance: On 12Jun09 using eight streams between two A+ systems connected via eight 10GE's, measured an aggregate performance of 69.2907 Gbps unidirectional, and bidirectional 38.6955 Gbps transmit & 38.5842 Gbps receive (77.2797 Gbps total aggregate)
- "A-" systems:
  - Primarily for network throughput evaluations via nuttcp memory-to-memory testing at up to 20-Gbps unidirectional, 20-Gbps bidirectional (40-Gbps "total")

Introduction To GSFC High End Computing 20, 40 &100 Gbps Network Testbeds

Phase 1.1 Network-Test Workstation Functional Objectives with Performance Targets

- "A++" systems:
  - Primarily for network throughput evaluations via nuttcp memory-to-memory testing at up to 100-Gbps unidirectional, 50-Gbps bidirectional (100-Gbps "total")
- Actual performance "in-progress"
  - On 6Aug09 measured an aggregate performance of 100.4637 Gbps in transmits; but currently only up to 56.4703 Gbps in receives
    - Test configuration has each of the two quad-core Xeon processors of one A++ system connected via six 10GE's to one of two quad-core i7-based A+ systems
    - Twelve streams are generated – one for each of the twelve 10GE connections handled by the one A++ system

Introduction To GSFC High End Computing 20, 40 &100 Gbps Network Testbeds

Approximate Costs (With components acquired via SEWP IV in lot-sizes of 3 - 15, and self assembly) of Phase 1 & 1.1 Network-Test Workstations

- "B" System: ~$6.8K
- "C" System: ~$9.0K
- "A" System: ~$4.6K
- "A+" System: ~$6.5K
- "A-" System: ~$3.6K
- "A++" System: ~$11.1K
- For more detail, contact Paul.Lang@nasa.gov

Overall Timeline for the HEC 20, 40 &100 Gbps Network Testbeds

```
                  | FY10   | FY11   | FY12   | FY13   | FY14   | FY15
4x10G
A. LAN+MAN
   Testing ttttttttttttttt
B. WAN-to-anywhere
   Testing   ttttttttttttttt

1x40G
A. LAN+MAN
   Testing       ttttttttooooooooooooooo
B. WAN betw ARC+GSFC
   Testing               TTTTTTTTTTTTTTT
C. WAN betw ARC+GSFC
   Operations                    OOOOOOOOOOOOOOO

1x100G
A. LAN+MAN
   Testing                       ttttttttooooooooooooooooo
B. WAN betw ARC+GSFC
   Testing                                TTTTTTTTTTTTTTT
C. WAN betw ARC+GSFC
   Operations                                    OOOOOOOOO
```

Wherein:

ttt…ttt implies only pre-operational testing use

ooo…ooo implies operational use in LAN+MAN

TTT…TTT implies only pre-operational testing use in WAN

OOO…OOO implies operational use in WAN

Introduction To GSFC High End Computing 20, 40 &100 Gbps Network Testbeds

Summary of HEC Test WAN Requirements

- Key SLA Requirements of HEC's WAN Test Network
  - Availability (percent): 80.0
  - Restoral Time: <48 hours
  - Coverage Period: 24x7
  - RTT between sites in CONUS: <100ms
  - Packet Loss: <1E-7
  - Jumbo Frames: Transport of 9000-byte IP MTU jumbo frames without fragmentation
  - Bandwidth between NAS@ARC and NCCS@GSFC:
    - 1Jul11 through 30Jun13: 40Gbps (i.e., 1x40G, not 4x10G or other LAG approaches)
    - 1Jul13 through 30Jun15: 100Gbps (i.e., 1x100G, not 10x10G or other LAG approaches)
- The above SLA requirements apply only to the "TTT…TTT" links
- The SLA requirements of "OOO…OOO" links are similar to "TTT…TTT" links except that "OOO…OOO" links have their Availability (percent) parameters at 99.50 and their Restoral Time parameters at 4 hours; and therefore they likely need to be provisioned from a different supplier than "TTT…TTT" links

Introduction To
GSFC High End Computing
20, 40 &100 Gbps Network Testbeds

Current Candidate 40Gbps MAN & WAN Pathways For Use During Early Stages Of Phase 1 20Gbps & Phase 2 40Gbps Testbeds

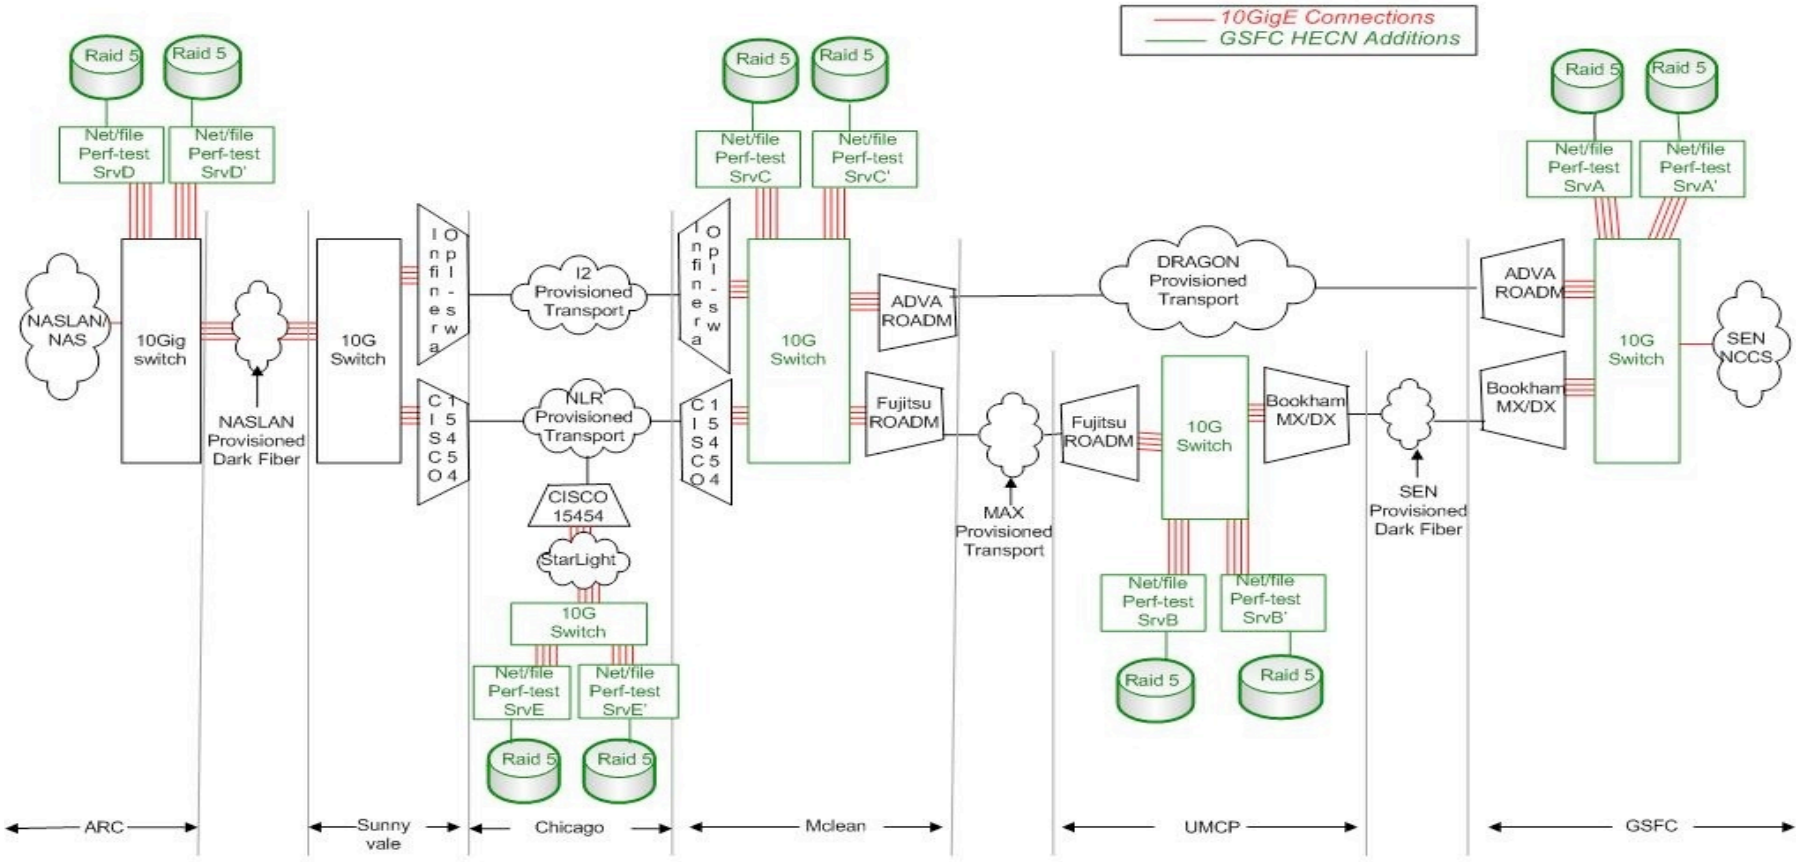

J.P.Gary/A.Muppalla
6/17/09

Introduction To GSFC High End Computing 20, 40 &100 Gbps Network Testbeds

Key Notes Regarding the Previous Diagram (1 of 2)

- The diagram's title is "Current Candidate 40-Gbps MAN & WAN Pathways For Use During Early Stages of Phase 1 20-Gbps and Phase 2 40-Gbps Testbeds"; and the referenced Testbeds were identified and/or described in [1] and [2]
- Where different but parallel pathways are shown (such as Internet2 vs NLR, and DRAGON's ADVA vs MAX's Fujitsu MAN pathways to MCLN) only one WAN and one MAN pathway is needed; but we'll have enough ports on our 10-GE switches to accommodate both WAN pathways and both MAN pathways without re-cabling if the opportunities should arise
- In the Early Stages, 4x10-GE link aggregation (LAG) is key (vs later stages if/when 40-GE Media Access Controllers (MACs) become available and affordable)

Introduction To GSFC High End Computing 20, 40 &100 Gbps Network Testbeds

Key Notes Regarding the Previous Diagram (2 of 2)

- Use of Layer 1+2 DCN-enabled VLANs versus Layer 1+2+3 full IP routed networks in both the Regional/MAN and WAN testbeds is critical
  - Sufficient to enable more effort to be focused on the primary subjects of this plan/proposal which are the processor interfaces and LAN infrastructure needed on the ends of the intervening links
  - Core IP routing issues (while otherwise interesting with many R&D challenges remaining) are not the primary subject of this plan/proposal
  - Costs of 40 and 100 Gbps Layer 1+2+3 router interfaces are likely to be two or more orders of magnitude greater than 40 and 100 Gbps Layer 1+2 Ethernet switch interfaces which are sufficient to enable the needed VLANs

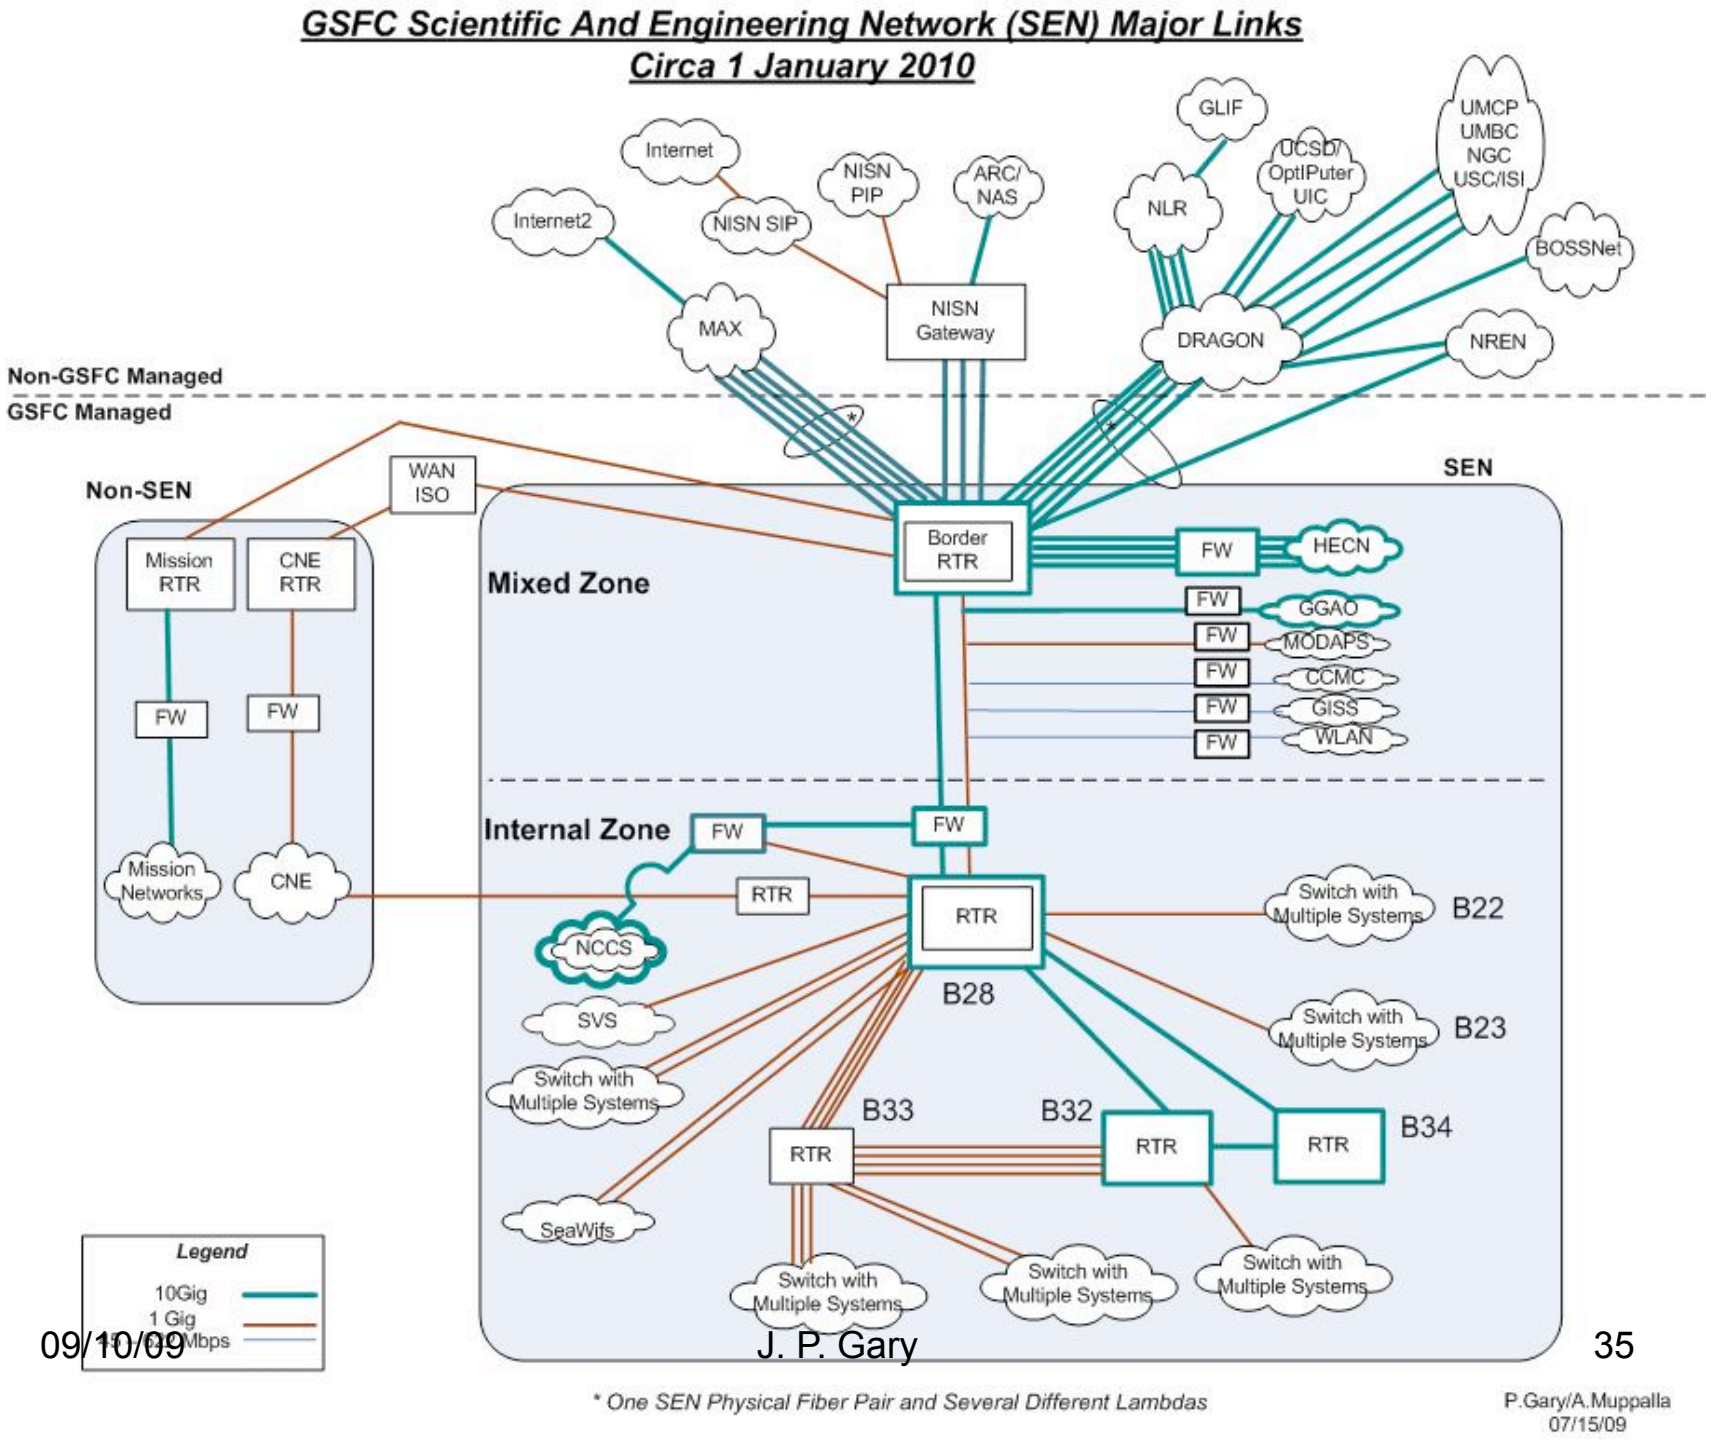
GSFC Scientific And Engineering Network (SEN) Major Links
Circa 1 January 2010
Internet
Internet2
NISN SIP
NISN PIP
ARC/ NAS
NISN Gateway
MAX
GLIF
NLR
UCSD/ OptIPuter UIC
UMCP UMBC NGC USC/ISI
BOSSNet
DRAGON
NREN
Non-GSFC Managed
GSFC Managed
Non-SEN
WAN ISO
SEN
Mission RTR
CNE RTR
FW
Mission Networks
CNE
Border RTR
Mixed Zone
HECN
GGAO
MODAPS
CCMC
GISS
WLAN
Internal Zone
RTR
NCCS
SVS
Switch with Multiple Systems
SeaWifs
B28
B22
B23
B33
B32
B34
Legend
10Gig
1 Gig
* One SEN Physical Fiber Pair and Several Different Lambdas
P.Gary/A.Muppalla
07/15/09

Introduction To GSFC High End Computing 20, 40 &100 Gbps Network Testbeds

Next Steps (Near Term) with Key Challenges

- Phase 1 network-test workstation assembly, bench tests, and deployments - LAN, MAN, WAN
  - Inclusion of application testing, as scripts now are only partially complete
  - Adequacy of disk subsystem and application tuning to achieve throughput performance objectives
- Investigation of candidate Phase 2 network-test workstations, refinement of Phase 2 throughput performance objectives, and refinement of Phase 2 LAN, MAN, WAN pathway targets
  - Opportunities to obtain early 40 & 100 GE-MAC-interfaces in the Phase 2 testbeds
- Leveraging NASA's "in-formation" Emerging Network Technology Testbed initiative

Source: Jim Fischer/GSFC 606.0

"Emerging Network Technology" Testbed

High-end network technology collaboration within NASA

Scope:

- ✓Prototype advanced commercial networking capabilities in collaboration with stakeholders
- ✓Maintain NASA's dedicated advanced peering capability
- ✓Support integration testing infrastructure
- ✓Assure Civil Servant high-end networking technical expertise

In Formation by: OCIO NISN and HEC

Emerging Network Technology Testbed

Staffed by ARC, GSFC, MSFC

Managed Relationships

Including annual negotiation of experiments

Stakeholders

- NASA Integrated Services Network (NISN) — *Led by MSFC*
- High-End Computing Capability (HECC) — *HEC Facility at ARC*
- NASA Center for Computational Sciences (NCCS) — *HEC Facility at GSFC*
- ...


NETWORK
BOTTLENECKS

Introduction To
GSFC High End Computing
20, 40 &100 Gbps Network Testbeds

Q & A

GSFC SEN+HECN Summary Information

Backup Slides

GSFC SEN+HECN Summary Information

Some Key Data Flow/Path Restrictions Affecting CNE & SEN

[By GSFC security policy and/or management agreement]

- Only CNE-sourced traffic with SEN-destinations
  or SEN-sourced traffic with CNE-destinations
  can flow across the "Only intra-CNE-SEN routing" path between CNE & SEN

- CNE-source traffic can **not** traverse the SEN or its Border Router/FW to other networks; SEN-source traffic can **not** traverse the CNE or its Border Router/FW to other networks
  - Rationale:
    - The security policies/rulesets of the respective Border Routers/FWs intentionally block data flows transiting to/from other networks
    - The respective CNE & SEN Helpdesks are not adequately trained or equipped to support data flows transiting to/from other networks

The Only Flows Permitted Across Intra-CNE-SEN Router By GSFC Security Policies and/or Management Agreements

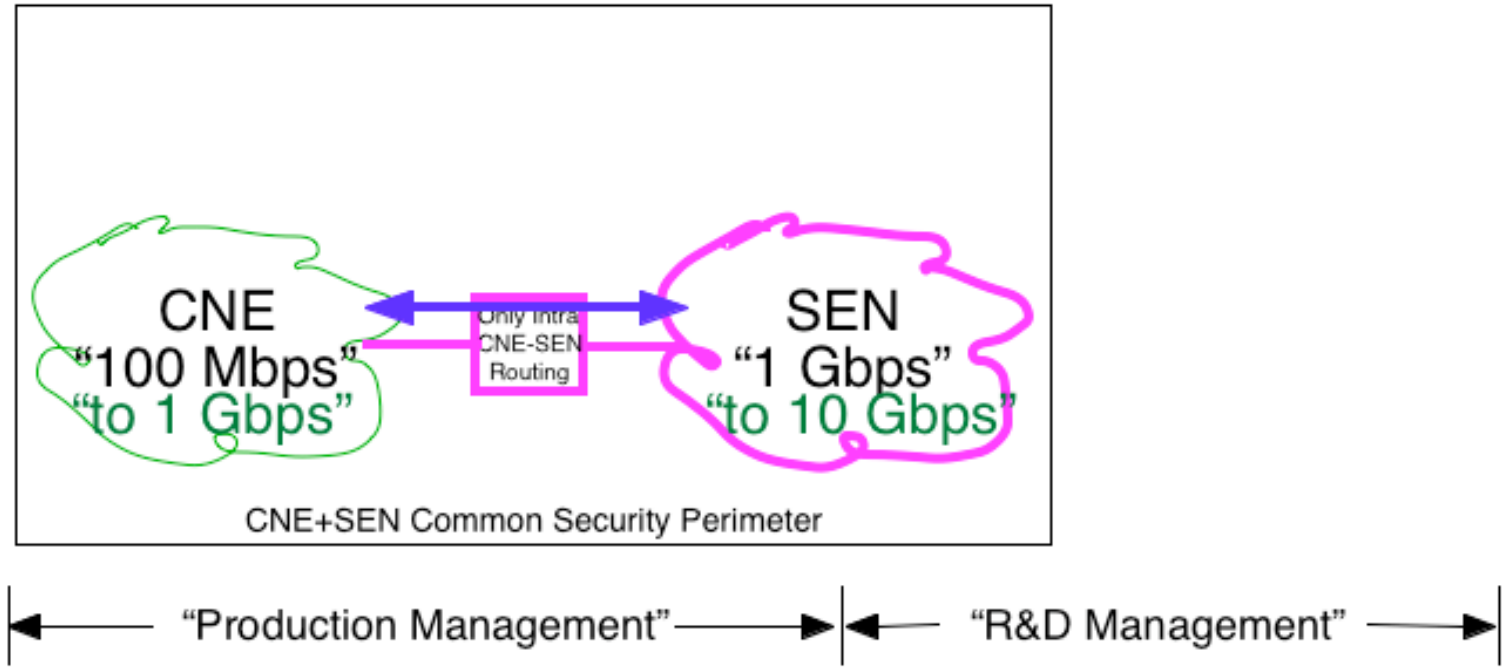

Examples of Traversal Flows Not Permitted By GSFC Security Policies and/or Management Agreements

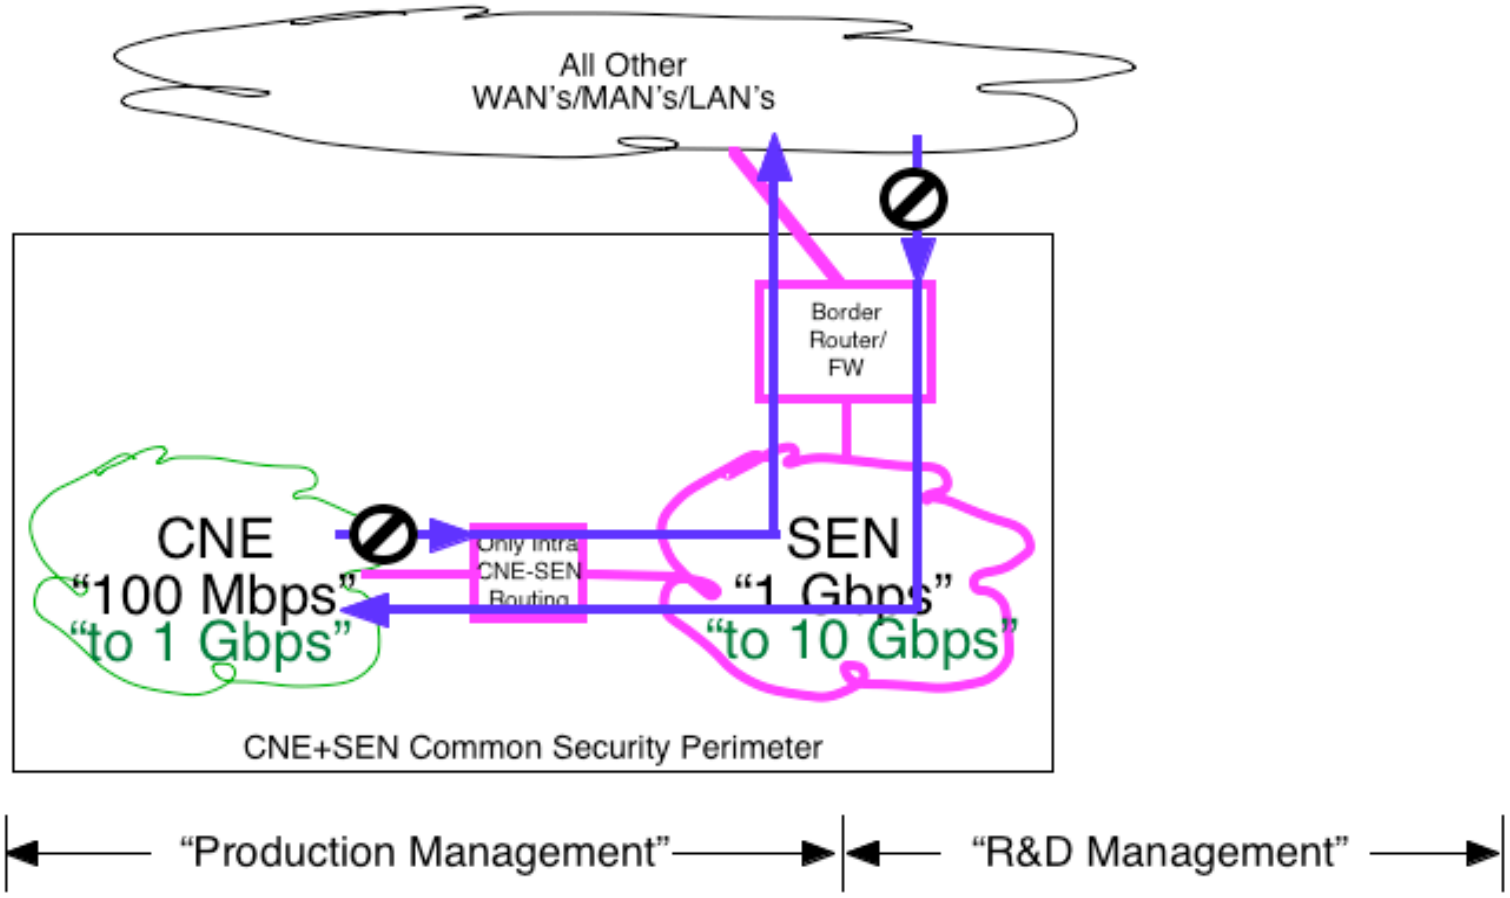

Introduction To GSFC High End Computing 20, 40 &100 Gbps Network Testbeds

Nuttcp (pronounced as new-t-t-c-p or nut-t-c-p)

- Primary author Bill Fink (william.e.fink@nasa.gov), with Rob Scott (rob@hpcmo.hpc.mil).
- Great follow-on to original ttcp network throughput performance measurement and troubleshooting tool. Highly recommended. One of the best!
- Over 60 examples of use included in Phil Dykstra's noteworthy tutorial for High Performance Data Transfer (at SC0x's).
- Advanced capabilities/features/options still being added to enable more sophisticated use, while retaining ease-of-use defaults.
- At http://www.nuttcp.net & included in perfSONAR's liveCD.

Introduction To GSFC High End Computing 20, 40 &100 Gbps Network Testbeds

40 Gbps Network Testing Between CLPK and MCLN

- Very successful; very informative
- Summary in "Industry Leaders Collaborate on 40 Gbps Live Network Trial" press release at http://www.businesswire.com/portal/site/google/?ndmViewId=news_view&newsId=20081110005564&newsLang=en

BUT

- No file copying application testing
- "Toothpaste-looping" flows, while innovative and sufficient for intervening network testing, are inappropriate for file copy testing
- Need higher performing network-test workstations for investigations focused on application throughput limitations

Introduction To GSFC High End Computing 20, 40 &100 Gbps Network Testbeds

Even 100 Gbps Networks Are Beginning To Emerge

- The Internet2 and the National LambdaRail (NLR) are upgrading their infrastructure to support 40-to-100 Gbps wavelengths
- IEEE P802.3ba Task Force has planned its Working Group Balloting starting in March 2009 and full completion of both the 40 and 100 Gbps Ethernet (GE) standards by mid-2010 (http://www.ieee802.org/3/ba/index.html)
- A partnership among ESNet, Internet2, Juniper Networks, Infinera, and Level3 Communications has announced the formation of a 100 GE testbed (https://mail.internet2.edu/wws/arc/i2-news/2008-11/msg00000.html)
- http://www.hpcwire.com/offthewire/ESnet-Receives-62M-to-Develop-Worlds-Fastest-Computer-Network-52989552.html?viewAll=y

IEEE P802.3ba Task Force Timeline

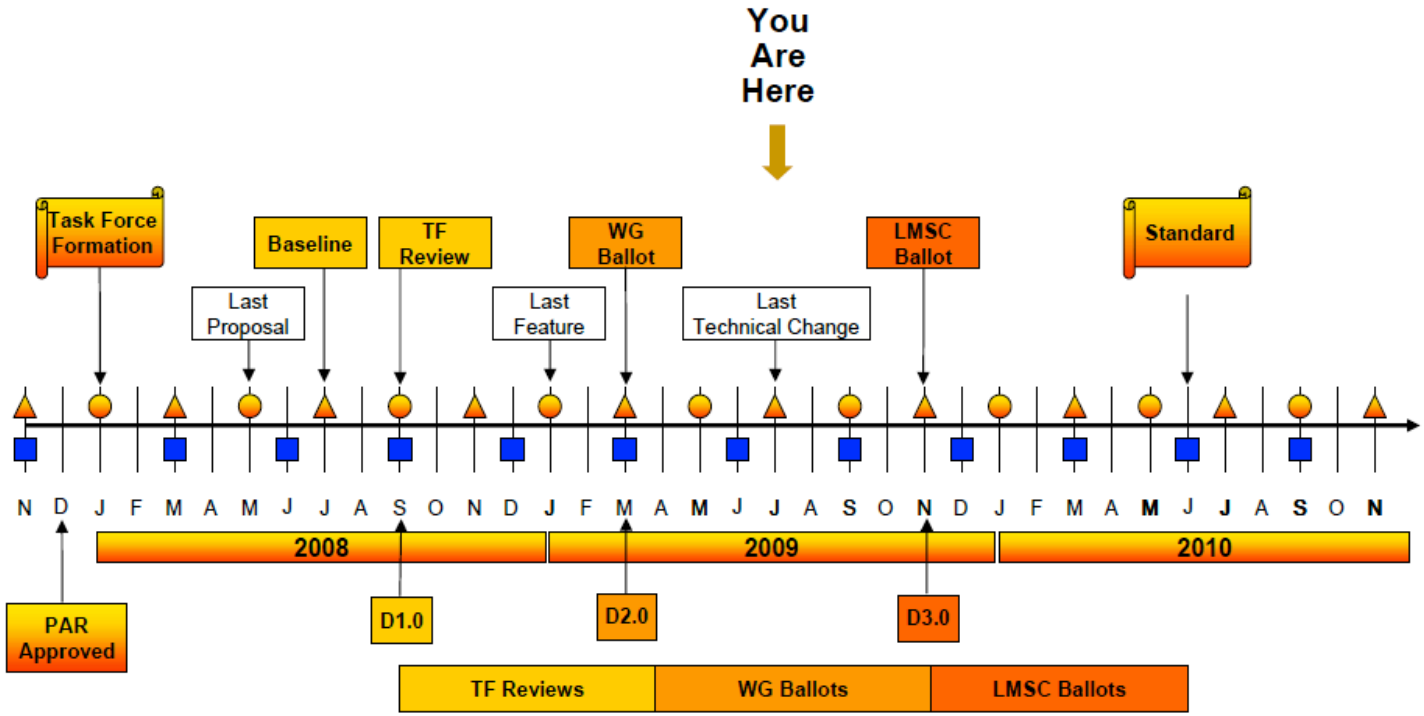

* Adopted by IEEE P802.3ba TF at March 08 Plenary

Legend
- IEEE 802 Plenary
- IEEE 802.3 Interim
- IEEE-SA Standards Board

Representative Architecture

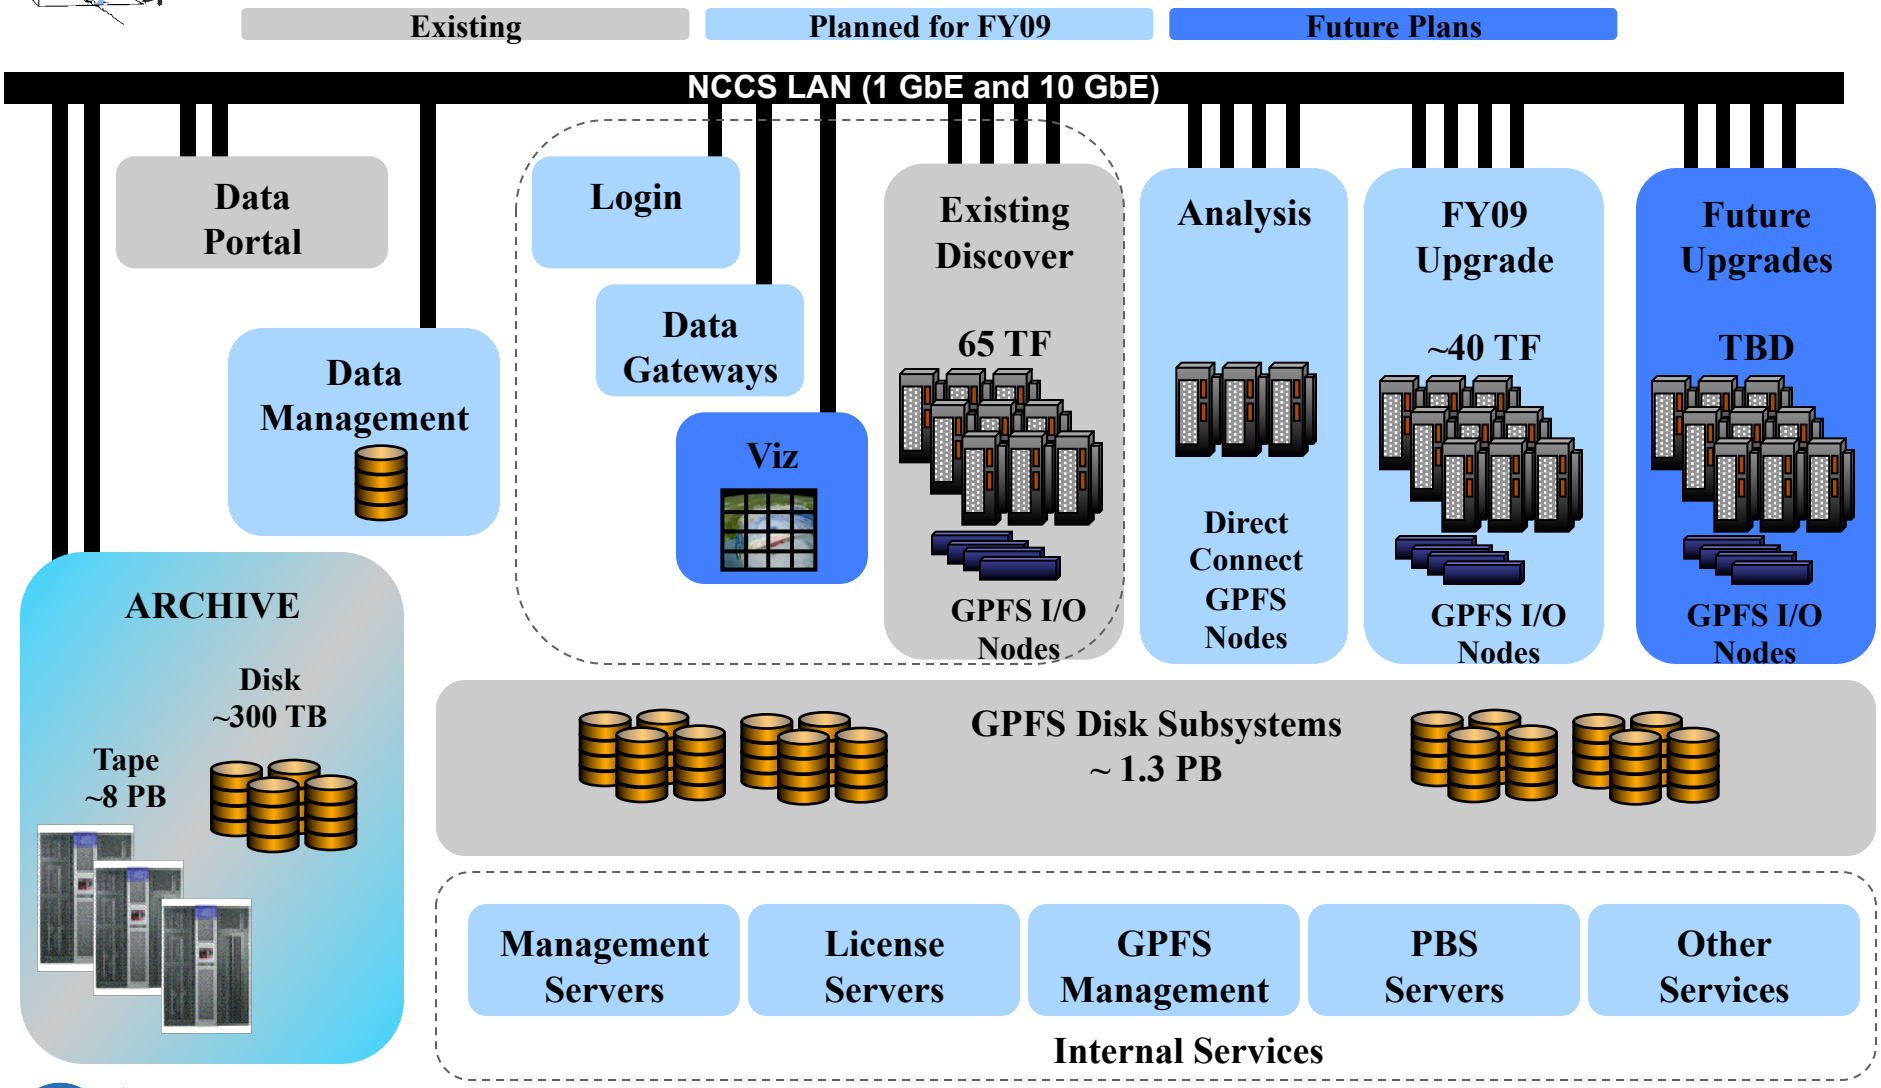

Excerpt from iRODS Update 2009 03 15.ppt
Sources: Dan Duffy/GSFC 606.2(NCCS) & Hoot Thompson/PTP(NCCS)

Integrated Rule-Oriented Data System (iRODS)

•Data grid software system developed by the Data Intensive Cyber Environments (DICE) group (developers of the SRB, the Storage Resource Broker), and collaborators.

Basic iRODS Components

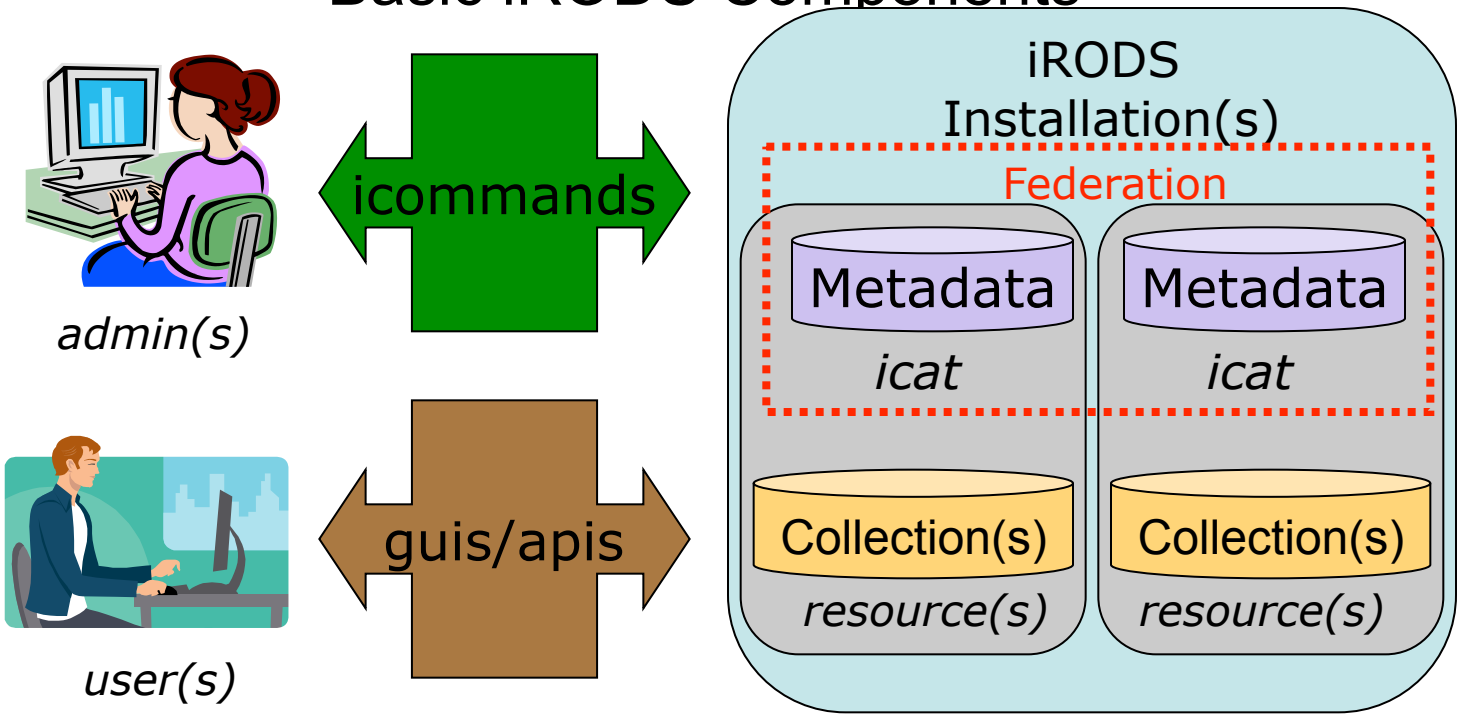

Introduction To GSFC High End Computing 20, 40 &100 Gbps Network Testbeds

Phase 1 Network-Test Workstations: Nominal "B" System

- Chassis: Supermicro 836TQ-R800B (3u 16bay 7slot 800W RPS)
- Motherboard: Asus P6T6 WS Revolution (5 PCIe V2 x8)
- Processors: one Intel i7 965 (3.2GHz quad-core Nehalem)
- Memory: Kingston KHX16000D3ULT1K3 (6GB 2000MHz DDR3 CL8)
- System disks: one Western Digital WD2500BEKT (2.5" 250GB)
- NICs: two Myricom 10G-PCIE2-8B2-2S+E (Dual 10GE SFP+)
- Raid controllers: two HighPoint RocketRaid 4320 (internal, 8 disks each)
- User disks: 16 Western Digital WD5001AALS (500GB)
- IB HCA: one Qlogic QLE7280 (DDR, 8x)
- For more detail, contact Paul.Lang@nasa.gov

Introduction To GSFC High End Computing 20, 40 &100 Gbps Network Testbeds

Phase 1 Network-Test Workstations: Nominal "C" System

- Nominal "B" (Baseline) System
- Minus:
  - NICs: One Myricom 10G-PCIE2-8B2-2S+E (Dual 10GE SFP+)
  - IB HCA: one Voltaire (DDR, 8x)
- Plus:
  - Raid controllers: two HighPoint RocketRaid 4322 (external, 8 disks each)
- Plus via SAS-connection:
  - Chassis: one Supermicro 836TQ-R800B (3u 16bay 7slot 800W RPS) with SAS converter/adaptor and cables
  - User disks: 16 Western Digital WD5001AALS (500GB)
- For more detail, contact Paul.Lang@nasa.gov

Introduction To GSFC High End Computing 20, 40 &100 Gbps Network Testbeds

Phase 1 Network-Test Workstations: Nominal "A" System

- Nominal "B" (Baseline) System
- Minus:
  - Raid controllers: two HighPoint RocketRaid 4320 (internal, 8 disks each)
  - User disks: 16 Western Digital WD5001AALS (500GB)
  - IB HCA: one Voltaire (DDR, 8x)

- For more detail, contact Paul.Lang@nasa.gov

Introduction To GSFC High End Computing 20, 40 &100 Gbps Network Testbeds

Phase 1 Network-Test Workstations: "A+" System

- Nominal "A" System
- Plus:
  - NICs: Two Myricom 10G-PCIE2-8B2-2S+E (Dual 10GE SFP+)

- For more detail, contact Paul.Lang@nasa.gov

Introduction To GSFC High End Computing 20, 40 &100 Gbps Network Testbeds

Phase 1 Network-Test Workstations: "A-" System

- Nominal "A" System
- Minus:
  - NICs: One Myricom 10G-PCIE2-8B2-2S+E (Dual 10GE SFP+)

- For more detail, contact Paul.Lang@nasa.gov

Introduction To GSFC High End Computing 20, 40 &100 Gbps Network Testbeds

Phase 1.1 Network-Test Workstations: "A++" System

- Chassis: Supermicro 836TQ-R800B (3u 16bay 7slot 800W RPS)
- Motherboard: Supermicro X8DAH+-F (6 PCIe V2 (4 x8 & 2 x16))
- Processors: two XEON W5580 (3.2GHz quad-core Nehalem)
- Memory: Kingston KHX16000D3ULT1K3 (6GB 2000MHz DDR3 CL8, running at 1333MHz)
- System disks: one Western Digital WD2500BEKT (2.5" 250GB)
- NICs: six Myricom 10G-PCIE2-8B2-2S+E (Dual 10GE SFP+)

- For more detail, contact Paul.Lang@nasa.gov

Introduction To GSFC High End Computing 20, 40 &100 Gbps Network Testbeds

Example Deployments

- "B" and "C" Systems will be deployed in the NCCS' GSFC-local "next generation" iRODS prototyping testbed
- "A" Systems will be deployed as firewalls in the GSFC's Science and Engineering Network (SEN)
- "A" Systems will be deployed as WAN delay emulators in the HECN Team's GSFC-local advanced networking testbed
- "A+" Systems will be deployed HECN Team's GSFC-local advanced networking testbed
- "A-" Systems may be deployed to JPL and LaRC for High End Computing Program testing across NISN's WANX upgrade
- For more detail, contact Pat.Gary@nasa.gov

Introduction To GSFC High End Computing 20, 40 &100 Gbps Network Testbeds

Notional Milestone Schedule

```
                          |    CY09    |    CY10     |    CY11    |
                          JFMAMJJASONDJFMAMJJASONDJFMAMJJASOND
Phase 1 20 Gbps Testbeds
O LAN & Region/MAN         ------******
O WAN                         ------******
Phase 2 40 Gbps Testbeds
O LAN & Region/MAN               ------******
O WAN                                ------******
Phase 3 100 Gbps Testbeds
O LAN & Region/MAN                  ------******
O WAN                                     ------******
```

Legend for Milestone Schedule

------ Planning and acquisition subphase

****** I&T plus demo subphase

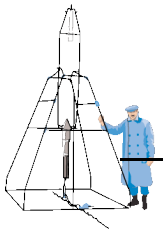

Plan/Proposal For Initiation of 20, 40 & 100 Gbps Network Technology Testbeds

"In Progress" Milestone Schedule

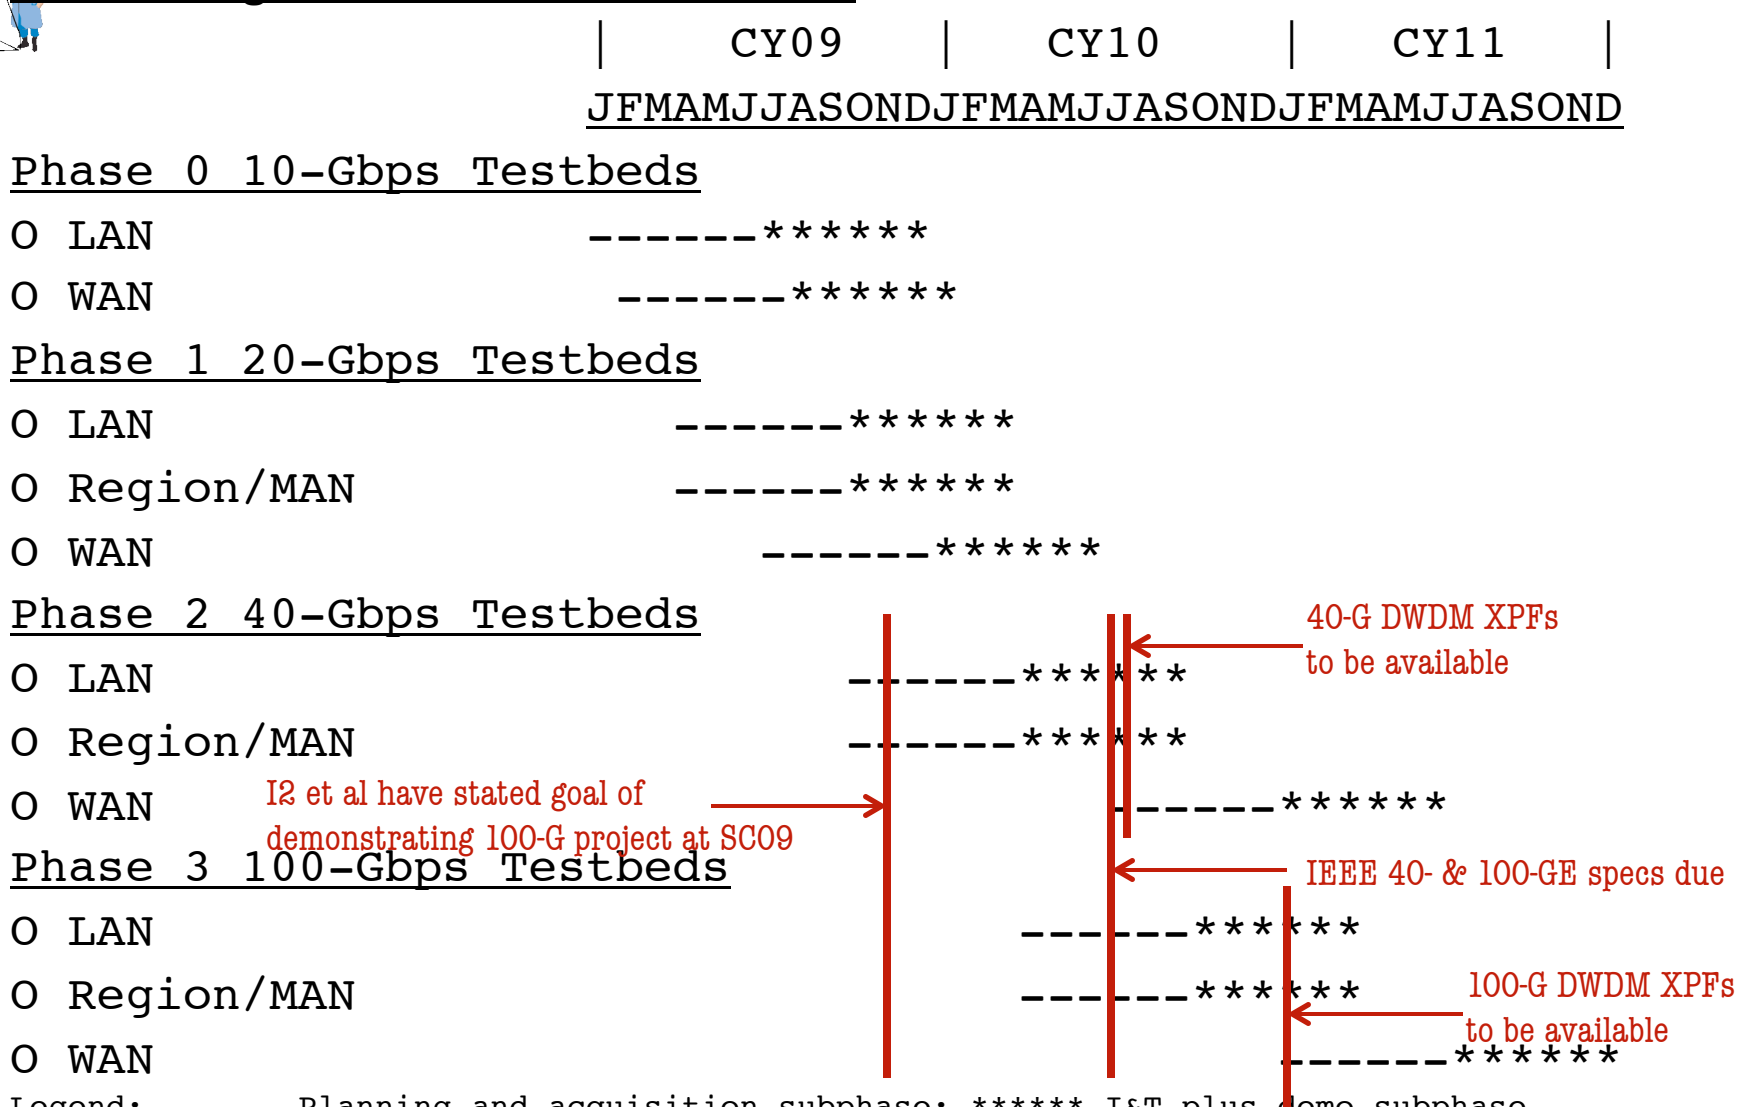

Introduction To GSFC High End Computing 20, 40 &100 Gbps Network Testbeds

Part of the HECN GSFC-local Advanced Networking Testbed

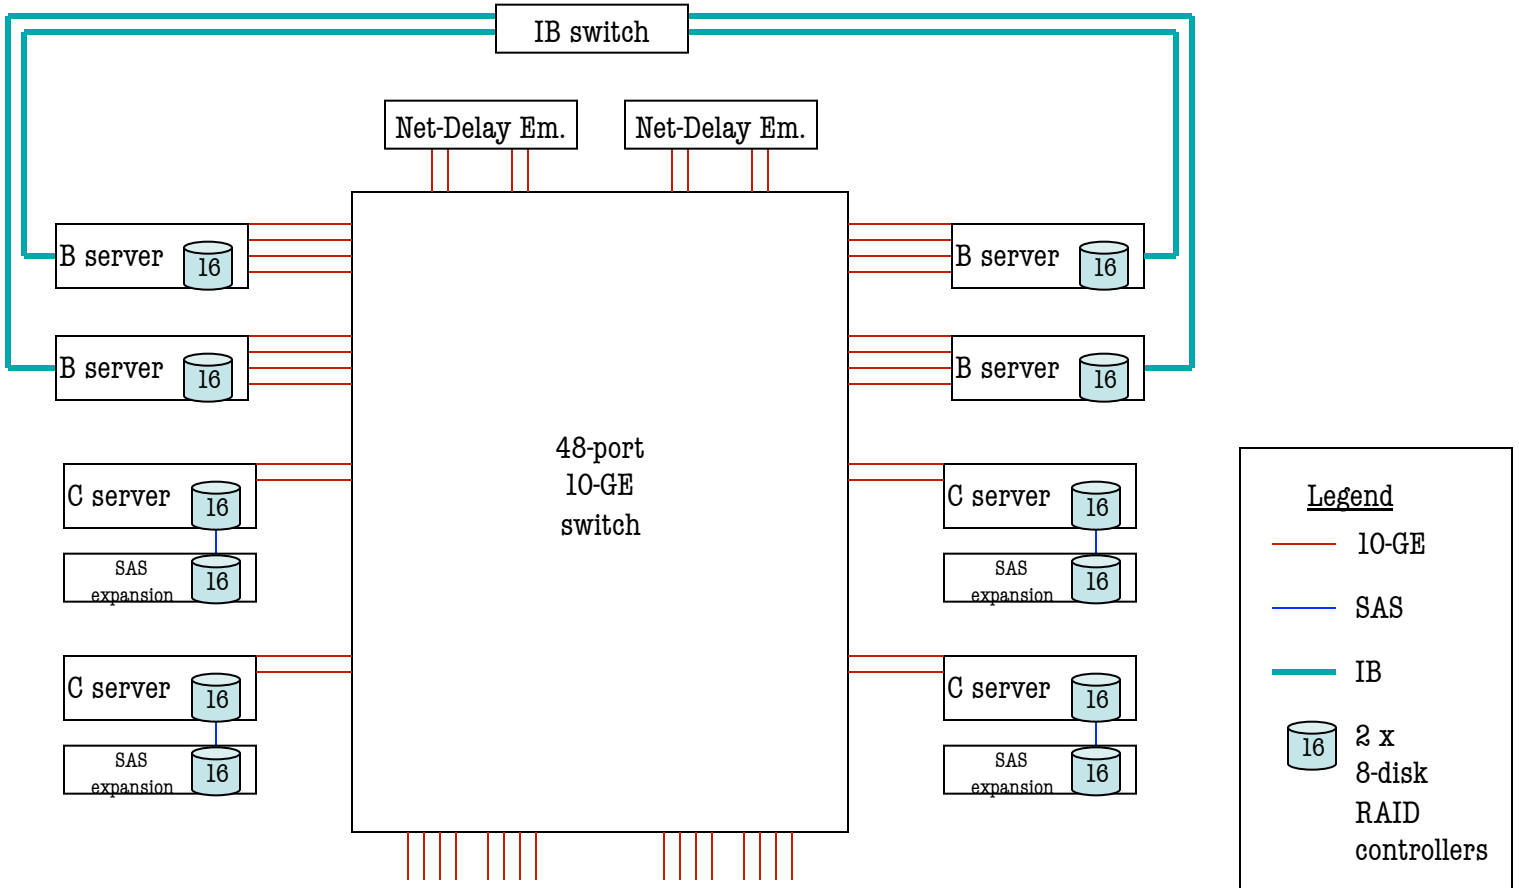

Introduction To GSFC High End Computing 20, 40 &100 Gbps Network Testbeds

Excerpt from iRODS Update 2009 03 15.ppt
Source: Dan Duffy/GSFC 606.2(NCCS) & Hoot Thompson/PTP(NCCS)

iRODS Prototype

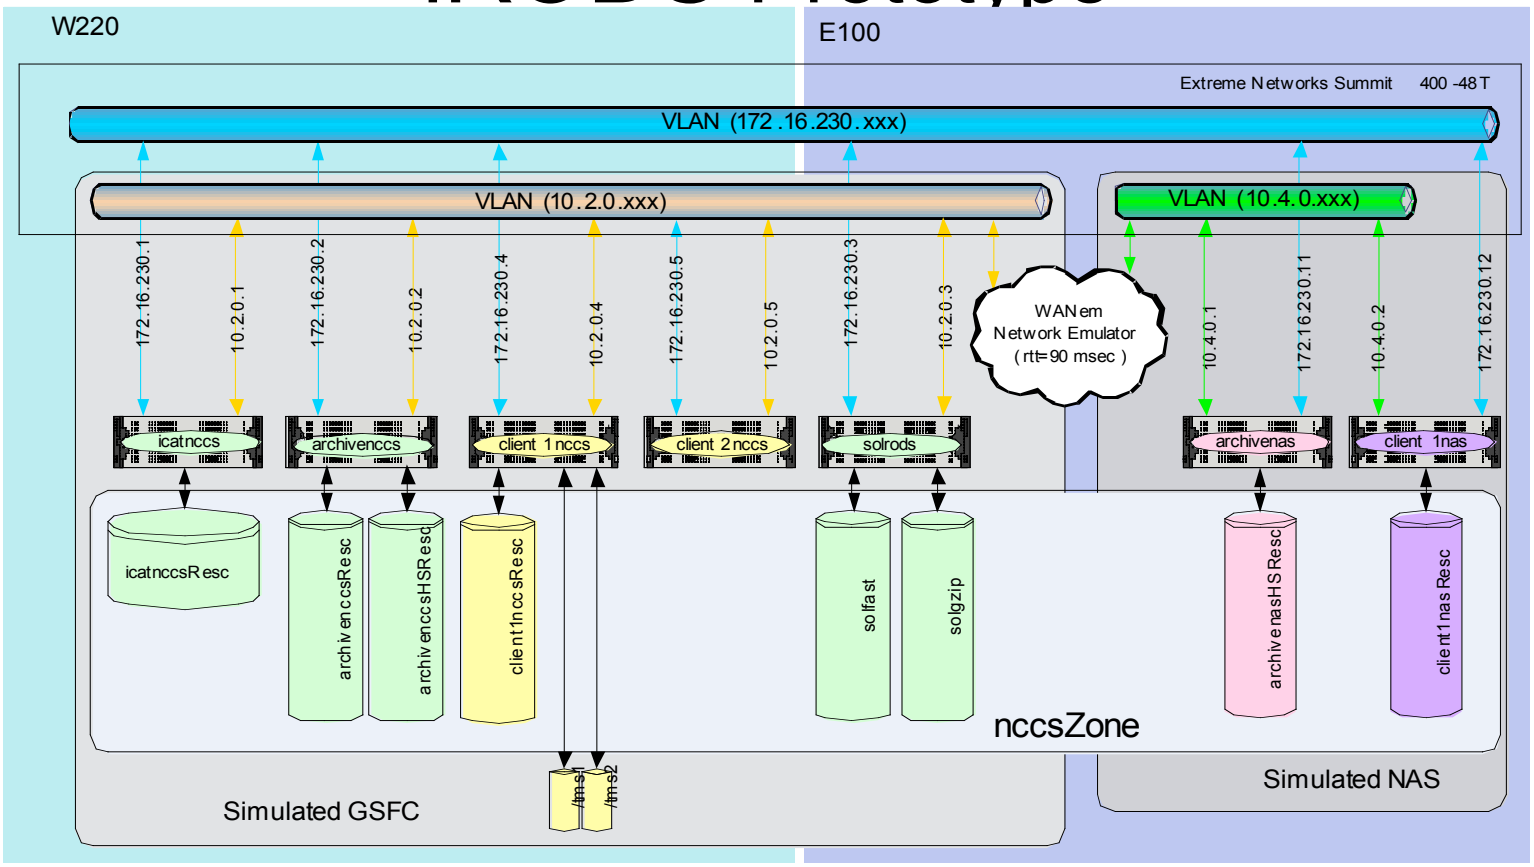

Introduction To
GSFC High End Computing
20, 40 &100 Gbps Network Testbeds

Candidate Use of the HECN GSFC-local Advanced Networking Testbed
For "Next Generation" NCCS iRODS Prototyping